



10:551 U.S. PTO
09/575145
11/5/23/111

Patent Office
Canberra

I, LEANNE MYNOTT, TEAM LEADER EXAMINATION SUPPORT AND SALES hereby certify that annexed is a true copy of the Provisional specification in connection with Application No. PQ 3632 for a patent by SILVERBROOK RESEARCH PTY LTD filed on 25 October 1999.

WITNESS my hand this
Tenth day of April 2000

LEANNE MYNOTT
TEAM LEADER EXAMINATION
SUPPORT AND SALES



CERTIFIED COPY OF
PRIORITY DOCUMENT

Best Available Copy

2000-01-01

AUSTRALIA

PATENTS ACT 1990

PROVISIONAL SPECIFICATION

FOR THE INVENTION ENTITLED:-

"PAPER BASED INFORMATION INTERFACE"

The invention is described in the following statement:-

The present inventions relate to a paper based information interface.

The inventions have been developed primarily for electronic commerce and will be described hereinafter with reference to that application. However, it will be appreciated that the inventions are not limited to this particular field of use and are also suitable for
5 non-business related information transfers such as that used by libraries, government organisations or households.

According to a first aspect of the invention there is provided a system for allowing a user to interact with one or more computers, the system including pages upon which are disposed tag elements, such that designation by the user of any one of the tag
10 elements results in instructions being applied in relation to the computer.

Preferably, some or all of the printed elements include identity data, the identity data for each tag element either:

uniquely identifying particular ones of the instructions generated in relation to the computer;

15 uniquely identifying a relative position of that tag element in relation to the page;
or

uniquely identifying a relative position of that tag element in relation to the page and uniquely identifying the page upon which that tag element is disposed.

Preferably, the identity data is not text or data that is otherwise semantically
20 understandable to a user.

Preferably, the user designates one of the tag elements by means of a tag element reader.

Preferably the reader takes the physical form of a stylus.

Preferably the pages are paper and the corresponding tag elements of each page are printed thereon.

In a preferred form of the first aspect, there is provided a system in which paper pages containing printed elements can allow a user to interact with a computer system
5 without having to explicitly enter control information into the computer system.

Preferably, the computer system is able to identify unique page elements, or unique page elements within unique pages, without explicit user knowledge.

Preferably, the preferred interaction technique uses an intelligent pen.

Preferably, the use of the term "paper" in these claims encompasses other media
10 such as plastic.

According to a second aspect of the invention there is provided a printer for printing pages, the pages being useable by a user to interact with a computer.

Preferably, the printer is configured to print one or more tag elements onto each of the pages, each of the tag elements including identity data, the identity data for each tag
15 element either:

uniquely identifying instructions useable for controlling the computer;

uniquely identifying a relative position of that tag element on its respective page;

or

uniquely identifying a relative position of that tag element on its respective page
20 and uniquely identifying the page upon which that tag element is disposed

Preferably, the printer is configured to print the tag elements using an ink that is invisible, or at least substantially not visible to the human eye. Preferably also, the

printer is configured to selectively print additional data onto the pages using one or more colours other than the invisible or substantially not visible ink.

In a preferred form of the second embodiment, there is provided a printer that can print documents as sets of pages of interactive paper.

5 According to a third aspect of the invention there is provided a printer configured to print a second document in response to selection of identity data disposed on a first document, wherein the identity data on the first document identifies the second document, thereby enabling retrieval of the second document from an electronic storage means for reproduction by the printer.

10 Preferably, each document is defined by a set of one or more associated pages, the printer being configured such that printing of a document includes printing all of the pages of that document.

Preferably, the identity data is read by a reader and transmitted to a receiver associated with the printer by a transmitter associated with the reader.

15 In a preferred form of the third embodiment, there is provided a printer that can print documents directly in response to requests initiated through interactive paper.

Preferably, the printer is configured to print documents as sets of pages of interactive paper.

20 According to a fourth aspect of the invention there is provided a system for allowing a user to interact with one or more computers, the system including pages upon which are disposed one or more text-based tag elements, such that designation by the user of any one of the tag elements results in instructions associated with that tag element being applied in relation to the computer.

Preferably, the instructions associated with the tag element cause a printer operatively connected with the computer to print out a page associated with the tag element.

Preferably, the tag element is linguistically or semantically understandable by at least some sub-group of users.

Preferably, the tag element takes the form of a hyperlink.

In a preferred embodiment of the fourth aspect, there is provided a system in which a word or phrase on interactive paper can cause, when selected by a user, a printer to print a document linked to that word or phrase.

According to a fifth aspect of the invention there is provided a system for allowing a user to interact with one or more computers, the system including pages upon which are disposed one or more graphical tag elements, such that designation by the user of any one of the tag elements results in instructions associated with that tag element being applied in relation to the computer.

Preferably, the instructions associated with the tag element cause a printer operatively connected with the computer to print out a page associated with the tag element.

Preferably, the tag element is linguistically or semantically understandable by at least some sub-group of users.

Preferably, the tag element takes the form of a hyperlink.

In a preferred embodiment of the fifth aspect, there is provided a system in which a graphic on interactive paper can cause, when selected by a user, a computer system to carry out an action linked to that graphic.

According to a sixth aspect of the invention there is provided a system of capturing stroke data in relation to one or more strokes generated by a user, the user applying the strokes by means of an active stylus to a page upon which is disposed a plurality of tag elements, each tag element including identity data defining a relative location of that tag
5 in relation to the page, the active stylus being configured to periodically read tags as each stroke is being generated, wherein the stroke data includes a sequential plurality of the identity data of tags that have been read by the stylus.

Preferably, the identity data is periodically read by the stylus and then transmitted to a memory module.

10 Preferably, the stylus includes a pressure sensor for detecting a pressure applied between the stylus and the page by a user. In this embodiment, it is preferable that the stroke data include pressure data. Preferably also, the pressure data is interspersed with the identity data or trans.

In other embodiments, at least a plurality of the identity data are timestamped.

15 In a preferred form of the sixth embodiment, there is provided a system in which freehand drawing by a user on interactive paper is captured by a computer system as "digital ink".

According to a seventh aspect of the invention there is provided a system of converting stroke data generated in accordance with the sixth aspect of the invention into
20 computer text data, the system including text recognition means for converting each stroke or combination of strokes into computer text data.

In a preferred form of the seventh aspect, there is provided a system in which text hand-written by a user on interactive paper is automatically converted by a computer system to computer text.

According to an eighth aspect of the invention there is provided a system of
5 verifying a signature or other handwritten authentication symbol on the basis of stroke data generated according to the sixth aspect, the system including means for processing the sequential plurality of identity data to verify that the signature falls within a predetermined range of geometrical tolerances.

Preferably, where the stylus includes a pressure sensor, verification of the
10 signature includes assessing whether the pressure associated with each of the sequential plurality of identity data is within a predetermined acceptable range.

Preferably, where the identity data is timestamped, verification of the signature includes assessing whether the total elapsed time implied by the timestamps, or relative times calculated between predetermined identity data of the plurality of sequential
15 identity, fall within a predetermined acceptable range.

Preferably, the system includes token comparison means for confirming that a token provided or entered into the system by the user corresponds with the signature. Preferably, the token takes the form of identification data provided by the stylus along with the stroke data, the identification data of the stylus being associated with one or
20 more potential signatures to be verified.

In a preferred form of the eighth aspect, there is provided a system in which a signature hand-written by a user on interactive paper can be automatically verified by a computer system.

Preferably, the signature is verified against known characteristics of the user's signature.

Preferably, the user's identity is provided explicitly to the computer system immediately before, during or immediately after generation of the signature for
5 verification via an identity token. Preferably, identity token is the user's intelligent pen.

According to a ninth aspect of the invention there is provided a digital printer including:

- (a) a printing mechanism for printing onto individual sheets of paper;
- (b) paper output means for outputting the sheets of paper after printing;
- 10 (c) a binding mechanism including:
 - (i) glue application means for applying glue to a predetermined portion of selected ones of the sheets of paper;
 - (ii) a collator for routing printed pages into contact with each other, such that those portions to which glue has been applied are in contact with appropriate portions on other
15 adjacent sheets of paper.

Preferably, the printer further includes pressure application means for squeezing the glued portions together with corresponding portions on the adjacent sheets of the paper.

Preferably, glue is applied to a single edge of a plurality of sheets of paper in a
20 bundle of sheets to be glued together.

Preferably, the digital printer is configured such that a multiple page document is recognised as such, and is glued together in a discrete bundle.

In a preferred form of the ninth aspect, there is provided a digital printer which automatically binds the pages of a document together using an internal binding mechanism, including one or more of a gluing mechanism, a stapling mechanism, or other chemical or mechanical binding means.

5 According to a tenth aspect of the invention there is provided a cartridge for use with the digital printer of the ninth aspect, the cartridge including:

 a supply of one or more inks or colorants for the printer; and

 a supply of adhesive for supply to the binding mechanism.

 According to an eleventh aspect of the invention there is provided a method of
10 distributing content to a plurality of users via a network document distribution system, the method including the steps of:

 defining one or more groups of users, each group having associated with it predetermined page layouts, content elements or content definitions;

 defining an individual profile for one or more of the users, each individual profile
15 including predetermined page layouts, content elements or content definitions;

 for each group, broadcasting content associated with that group to the individual users within that group; and

 for each individual, pointcasting content associated with the user's profile to the user.

20 Preferably, the method further includes the step, for each user, of compiling the broadcast and pointcast content into a single document.

 According to a twelfth aspect of the invention there is provided a network document distribution system in which unique page elements and/or layouts are

pointcast to individual receivers and shared page elements and/or layouts are multicast to groups of receivers.

According to a thirteenth aspect of the invention there is provided a system in which the logical content of a printed page is persistently archived in such a way that any
5 user interaction with the printed page at a later time can be interpreted with respect to the logical content of the page.

According to a fourteenth aspect of the invention there is provided a system in which a user can make an online payment exclusively by marking and/or signing paper pages which represent payment instructions.

10 Preferably, payment card information is linked to the user's identity on the network. More preferably, the user's signature and the user's identity token (for example the user's intelligent pen) together allow the computer system to authenticate the payment. Payment forms can include statements delivered regularly to a user's printer, as well as online order forms printed interactively.

15 According to a fifteenth aspect of the invention there is provided a system in which a user can make an online purchase exclusively by marking paper pages printed on demand.

According to a sixteenth aspect of the invention there is provided a system in which a user can make an online purchase exclusively by marking pre-printed paper
20 pages.

Preferably, the pre-printed pages include flyers, brochures and catalogues delivered by traditional means.

According to a seventeenth aspect of the invention there is provided a system in which a user can browse the World Wide Web exclusively by viewing and marking paper pages printed on demand.

Preferably, hyperlinks embedded in Web pages remain active when printed on
5 interactive paper, and when selected request the printer to print the linked pages. More preferably, because of the greater resolution of the printed page, Web pages intended for screen display are printed in two columns and at about six times the density of screen display.

According to an eighteenth aspect of the invention there is provided a system in
10 which advertising is automatically placed on the reverse side of an interactive paper sheet if the reverse side is otherwise empty.

Preferably, the minimum unit of output from an interactive printer is a two-sided sheet of paper. In many interaction sequences the second side of the page will be empty, and thus provides space for advertising. The advertising is non-intrusive, and although it
15 may not be seen immediately, it persists on the page and may be seen later.

According to a nineteenth aspect of the invention there is provided a system in which a user can browse a directory of information exclusively by viewing and marking paper pages printed on demand.

Preferably, the directories, such as of people, places, services, and the like, are
20 ubiquitous. Assuming 200 entries per printed sheet, the standard unit of interaction, a global phone book for six billion people requires only five index levels, and a user can find a specific entry by printing only five sheets of interactive paper.

According to a twentieth aspect of the invention there is provided a system in which a user can obtain printed product information on demand via a product's universal product code (UPC) bar code, via a bar code scanner communicating directly with a printer.

5 Preferably, product information in the form of usage instructions, warranty information, promotions, and the like, are linked directly to any product which bears a universal product code bar code. More preferably, a user scans the product's bar code to request that the information be printed on the user's printer. More preferably, if the product information is printed as interactive paper pages, then the user can navigate a
10 whole collection of product information starting with just a single page. Even more preferably, the bar code scanner is ideally integrated into the user's intelligent pen.

According to a twenty first aspect of the invention there is provided a system in which interactive content for both the World Wide Web and interactive paper can be generated from a single content description.

15 According to a twenty second aspect of the invention there is provided a application system which can present both a World Wide Web user interface and an interactive paper user interface.

According to a twenty third aspect of the invention there is provided a system for providing printed information to a user that is obtained from a remote source, the system
20 including:

a user printer module that is provided by a first party to the user for interfacing the user with the source, the module being responsive to the user requesting first information from the source for obtaining the first information and generating a first printed media

that displays to the user the first information together with second information derived from a second party;

identifier means for applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed
5 media that displays to the user third information; and

calculation means responsive to the module for the determining a financial debit that is applied by the first party against the second party.

Preferably, the system includes a plurality of modules associated with corresponding users and the debit is a predetermined function of the number of printed
10 media generated that display the second information. In other embodiments, the debit is a predetermined function of the number of printed media generated that display the third information. In still further embodiments the debit is a predetermined function of both the number of printed media generated that display the second information and the number of printed media generated that display the third information.

15 Preferably also, the second party is a supplier of goods and/or services and the second information or the third information allows the user to purchase those goods and/or services at a predetermined purchase price. Preferably also, the debit is a predetermined function of the purchase price. For example, in some embodiments the debit is a flat fee for each purchase while in other embodiments the debit is a fixed
20 proportion of the purchase price. In still further embodiments, the fixed proportion varies and is dependent upon the particular goods and/or services purchased by the user.

In a preferred form, there are a plurality of second parties and the quantum of the debit is contingent upon one or more of the following characteristics of the respective second party:

- geographic location;
- 5 • quantity of purchases completed in a given time period; and
- relative geographic location with respect to the user.

Preferably, the identifier is a tag. More preferably, the identifier is a plurality of tags spaced apart on the respective printed media.

In some embodiments the identifier is a character string that the user provides to
10 the module to obtain the third information. For example, the string may include a sequence of numbers that are manually keyed into the module. More preferably, however, the string is manually scanned by the user.

According to a twenty fourth aspect of the invention there is provided a method for providing printed information to a user that is obtained from a remote source, the method
15 including the steps of:

having a first party provide to a user a user printer module for interfacing the user with the source, the module being responsive to the user requesting first information from the source for obtaining the first information and generating a first printed media that displays to the user the first information together with second information derived
20 from a second party;

applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user third information; and

determining, in response to the module, a financial debit that is applied by the first party against the second party.

According to a twenty fifth aspect of the invention there is provided a system for providing printed information to a user that is obtained from a remote source, the system including:

processor means for providing first information derived from a first party;

a user printer module that is provided by a second party to the user for interfacing the user with the source, the module being responsive to the user requesting second information from the source for obtaining the second information and generating a first printed media that displays to the user the second information together with the first information;

identifier means for applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user third information; and

payment means for receiving a predetermined payment from the first party to the second party and for generating a disable signal if that payment is not made wherein the generation of the disable signal prevents the module obtaining the second information.

According to a twenty sixth aspect of the invention there is provided a method for providing printed information to a user that is obtained from a remote source, the method including the steps of:

providing first information derived from a first party;

interfacing the user with the source with a user printer module that is provided by a second party, the module being responsive to the user requesting second information

from the source for obtaining the second information and generating a first printed media that displays to the user the second information together with the first information;

applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that

5 displays to the user third information; and

receiving a predetermined payment from the first party to the second party and for generating a disable signal if that payment is not made wherein the generation of the disable signal prevents the module obtaining the second information.

According to a twenty seventh aspect of the invention there is provided a system
10 for providing to a user printed first information about a first party that is obtained from a remote source, the system including:

memory means at the remote source for storing the first information;

a user printer module that is provided by a second party to the user for interfacing the user with the source, the module being responsive to the user requesting the first
15 information for obtaining the first information from the memory means and generating a first printed media that displays to the user the first information;

identifier means for applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user second information; and

20 calculation means responsive to the module for determining a financial debit that is applied by the second party against the first party.

According to a twenty eighth aspect of the invention there is provided a method for providing to a user printed first information about a first party that is obtained from a remote source, the method including the steps of:

storing the first information in memory means at the remote source;

5 interfacing the user with the source with a user printer module that is provided by a second party to the user, the module being responsive to the user requesting the first information for obtaining the first information from the memory means and generating a first printed media that displays to the user the first information;

applying an identifier to the first printed media such that designation of the
10 identifier by the user results in the module generating a second printed media that displays to the user second information; and

being responsive to the module for determining a financial debit that is applied by the second party against the first party.

According to a twenty ninth aspect of the invention there is provided a system for
15 providing printed information derived from a first party to a user where the information is obtained from a remote source, the system including:

a user printer module that is provided by a second party to the user for interfacing the user with the source, the module being responsive to the user requesting the information from the source for obtaining the information and generating a first printed
20 media that displays the information to the user;

identifier means for applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user additional printed information; and

calculation means responsive to the module for the determining a financial debit that is applied by the second party against the first party.

According to a thirtieth aspect of the invention there is provided a method for providing printed information derived from a first party to a user where the information
5 is obtained from a remote source, the method including the steps of:

interfacing the user with the source with a user printer module that is provided by a second party, the module being responsive to the user requesting the information from the source for obtaining the information and generating a first printed media that displays the information to the user;

10 applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user additional printed information; and

being responsive to the module for the determining a financial debit that is applied by the second party against the first party.

15 These aspects, in the preferred embodiments of the invention, are directed toward allowing a provider of an interactive printer to receives a percentage of advertising fees, click-through fees, sales commissions and/or transaction fees earned directly through the printer. This allows a profitable role for the provider of interactive printers, and allows the provider to attract users by being able to partially or fully subsidise the capital and
20 running costs of each user's printer.

According to a thirty first aspect of the invention there is provided a system for providing a user with a printed publication from a first party, the system including:

a publication source for storing publication data representative of the publication,
the publication data including first information about a second party;

a client server for selectively accessing the publication data;

a user printer module for interfacing the user with the server, the module being
5 responsive to the user requesting the printed publication for obtaining the publication
data and reproducing the printed publication on a print media for viewing by the user;

identifier means for applying an identifier to the print media such that designation
of the identifier by the user results in the module generating a second print media that
displays to the user second information; and

10 calculation means responsive to the module for determining a financial debit that is
applied by the first party against the second party.

Preferably, the debit is a predetermined function of the number of reproductions of
the printed publication that are generated. In other embodiments, the second information
is associated with the second party and the debit is a predetermined function of the
15 number of second print media generated. In still further embodiments the debit is a
predetermined function of both the number of printed publications generated that display
the first information and the number of second print media generated that display the
second information.

Preferably also, the second party is a supplier of goods and/or services and the first
20 information or the second information allows the user to purchase those goods and/or
services at a predetermined purchase price. Preferably also, the debit is a predetermined
function of the purchase price.

According to a thirty second aspect of the invention there is provided a method for providing a user with a printed publication from a first party, the method including the steps of:

- storing publication data representative of the publication in a publication source,
- 5 the publication data including first information about a second party;
- providing a client server for selectively accessing the publication data;
- interfacing the user with the server with a user printer module, the module being responsive to the user requesting the printed publication for obtaining the publication data and reproducing the printed publication on a print media for viewing by the user;
- 10 applying an identifier to the print media such that designation of the identifier by the user results in the module generating a second print media that displays to the user second information; and
- being responsive to the module for determining a financial debit that is applied by the first party against the second party.

- 15 According to a thirty third aspect of the invention there is provided a system for providing a user with a printed publication from a first party, the system including:

- a publication source for storing publication data representative of the publication, the publication data including first information about a second party;
- a client server for selectively accessing the publication data;
- 20 a user printer module for interfacing the user with the server, the module being responsive to the user requesting the printed publication for obtaining the publication data and reproducing the printed publication on a print media for viewing by the user;

identifier means for applying an identifier to the print media such that designation of the identifier by the user results in the module generating a second print media that displays to the user second information; and

payment means for receiving a predetermined payment from the second party to the first party and for generating a disable signal if that payment is not made wherein the generation of the disable signal prevents the reproduction of the first information.

According to a thirty fourth aspect of the invention there is provided a method for providing a user with a printed publication from a first party, the method including the steps of:

10 storing publication data representative of the publication, the publication data including first information about a second party;

selectively accessing the publication data with a client server;

interfacing the user with the server with a user printer module, the module being responsive to the user requesting the printed publication for obtaining the publication data and reproducing the printed publication on a print media for viewing by the user;

15 applying an identifier to the print media such that designation of the identifier by the user results in the module generating a second print media that displays to the user second information;

receiving a predetermined payment from the second party to the first party; and

20 generating a disable signal if that payment is not made wherein the generation of the disable signal prevents the reproduction of the first information.

These aspects, in the preferred embodiments of the invention, are directed toward allowing a publisher who publishes on interactive paper to receive advertising fees

and/or click-through fees on all advertising placed within their publications, and/or sales commissions on all sales initiated through the advertising. This advantageously provides a profitable role for a publisher who publishes on interactive paper, and allows the publisher to attract readers by being able to provide them with subsidised or free
5 publications.

According to a thirty fifth aspect of the invention there is provided a system for providing printed information to a user that is obtained from a data base of a first party, the system including:

server means for supporting the data base and storing second information derived
10 from a second party;

a user printer module for interfacing the user with the server means, the module being responsive to the user requesting first information from the data base for obtaining the first information and generating a first printed media that displays to the user the first information together with the second information;

15 identifier means for applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user third information; and

calculation means responsive to the module for the determining a financial debit that is applied by the first party against the second party.

20 According to a thirty sixth aspect of the invention there is provided a method for providing printed information to a user that is obtained from a remote source of a first party, the method including the steps of:

supporting the data base on server means which also stores second information derived from a second party;

interfacing the user with the server means with a user printer module, the module being responsive to the user requesting first information from the data base for obtaining
5 the first information and generating a first printed media that displays to the user the first information together with the second information;

identifier means for applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user third information; and

10 calculation means responsive to the module for the determining a financial debit that is applied by the first party against the second party.

According to a thirty seventh aspect of the invention there is provided a system for providing printed information to a user that is obtained from a data base of a first party, the system including:

15 server means for supporting the data base and for storing first information derived from a second party;

a user printer module that is provided by a second party to the user for interfacing the user with the server means, the module being responsive to the user requesting second information from the data base for obtaining the second information and
20 generating a first printed media that displays to the user the second information together with the first information;

identifier means for applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user third information; and

payment means for receiving a predetermined payment from the first party to the
5 second party and for generating a disable signal if that payment is not made wherein the generation of the disable signal prevents the module obtaining the second information.

According to a thirty eighth aspect of the invention there is provided a method for providing printed information to a user that is obtained from a data base of a first party, the method including the steps of:

10 supporting the data base on server means which also stores first information derived from a second party;

interfacing the user with the server means with a user printer module that is provided by a second party, the module being responsive to the user requesting second information from the data base for obtaining the second information and generating a
15 first printed media that displays to the user the second information together with the first information;

applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user third information; and

20 providing payment means for receiving a predetermined payment from the first party to the second party and for generating a disable signal if that payment is not made wherein the generation of the disable signal prevents the module obtaining the second information.

These aspects, in the preferred embodiments of the invention, are directed toward allowing a provider of persistent storage for the logical content of interactive paper pages to be paid out of a global account into which a percentage of all advertising fees, click-through fees, sales commissions and/or transaction fees earned on printers is placed.

- 5 This provides a profitable role for a persistent storage provider and allows the provision of page persistence to be de-coupled from page publication.

According to a thirty ninth aspect of the invention there is provided a system in which an account is maintained for an interactive printer.

- Preferably, the account receives credits for advertising printed on the printer,
10 debits for the amortised capital costs and running costs of the printer, and the user of the printer is billed on a regular basis for any debit balance. This allows a provider to subsidise the capital and running costs of a user's printer on the basis of advertising fees earned, while ensuring that no printer becomes a cost burden.

- According to a fortieth aspect of the invention there is provided a system in which
15 the balance of an interactive printer account is presented to a user as a collection of points which, if the balance is in credit, the user redeems for special products and/or services.

This provides an incentive for a user to print advertising, without necessarily giving the user an incentive to defraud the system.

- 20 According to a forty first aspect of the invention there is provided a system in which banking services are delivered via interactive paper.

Preferably, the banking services include support for balance enquiries, account transfers, bill payments and loan applications.

According to a forty second aspect of the invention there is provided a system in which insurance services are delivered via interactive paper.

Preferably, insurance services include support for policy negotiation, policy issue and claims.

5 According to a forty third aspect of the invention there is provided a business model in which travel services are delivered via interactive paper.

Preferably, travel services include support for enquiries about and reservations of flights, accommodation and holiday packages.

10 According to a forty fourth aspect of the invention there is provided a system in which a publisher sells the right to be linked to one of a set of generic category buttons prominently featured in a periodical publication delivered on interactive paper.

15 According to a forty fifth aspect of the invention there is provided a system for providing printed information to a user that is obtained from a data base of a first party, the data base having first information derived from a second party wherein the system includes:

a user printer module that is provided by a third party to the user for interfacing the user with the data base, the module being responsive to the user requesting the first information from the data base for obtaining the first information and generating a first printed media that displays to the user the first information;

20 identifier means for applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that displays to the user second information; and

calculation means responsive to the module for the determining a financial debit that is applied against the second party by the first party or the third party.

Preferably, the calculation means determines a first and a second financial debit that are applied against the second party by the first and third parties respectively.

5 According to a forty sixth aspect of the invention there is provided a method for providing printed information to a user that is obtained from a data base of a first party, the data base having first information derived from a second party wherein the method includes the steps of:

interfacing the user with the data base with a user printer module that is provided
10 by a third party, the module being responsive to the user requesting the first information from the data base for obtaining the first information and generating a first printed media that displays to the user the first information;

applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that
15 displays to the user second information; and

being responsive to the module for the determining a financial debit that is applied against the second party by the first party or the third party.

According to a forty seventh aspect of the invention there is provided a system for providing printed information to a user that is obtained from a data base of a first party,
20 the data base having first information derived from a second party wherein the system includes:

a user printer module that is provided by a third party to the user for interfacing the user with the data base, the module being responsive to the user requesting second

information from the data base for obtaining the second information and generating a first printed media that displays to the user the second information together with the first information;

identifier means for applying an identifier to the first printed media such that
5 designation of the identifier by the user results in the module generating a second printed media that displays to the user third information; and

calculation means responsive to the module for the determining a financial debit that is applied against the second party by the first party or the third party.

According to a forty eighth aspect of the invention there is provided a method for
10 providing printed information to a user that is obtained from a data base of a first party, the data base having first information derived from a second party wherein the method includes the steps of:

interfacing the user with the data base with a user printer module that is provided by a third party, the module being responsive to the user requesting second information
15 from the data base for obtaining the second information and generating a first printed media that displays to the user the second information together with the first information;

applying an identifier to the first printed media such that designation of the identifier by the user results in the module generating a second printed media that
20 displays to the user third information; and

being responsive to the module for the determining a financial debit that is applied against the second party by the first party or the third party.

Preferred embodiments of the above aspects of the invention, together with other aspects, will now be described, by way of example only, with reference to the following Appendices which are identified as:

Netpage, System Overview, draft version 0.4;

5 Netpage Printer, Design Description, draft version 0.3;

Netpage Pen, Design Description, draft version 0.3;

Netpage, Applications, draft version 0.2;

Netpage, Extensions, draft version 0.1; and

Self-Mapping Surfaces, Design Description, draft version 0.2.

10 Although the inventions have been described with reference to specific examples it will be appreciated by those skilled in the art that they may be embodied in many other forms.

DATED this 25th day of October, 1999


SILVERBROOK RESEARCH PTY LTD

15 Attorney: JOHN B. REDFERN
Fellow Institute of Patent Attorneys of Australia
of BALDWIN SHELSTON WATERS



Netpage System Overview

draft version 0.4, 20 October 1999



Silverbrook Research Pty Ltd
393 Darling Street, Balmain
NSW 2041 Australia
Phone: +61 2 9818 6633
Fax: +61 2 9818 6711
Email: info@silverbrook.com.au

Confidential

Contents

OVERVIEW	1
1 Introduction	2
1.1 Publishing to Netpage Printers	3
2 The Demise of Paper	5
3 News and Advertising Trends.....	7
3.1 Newspapers.....	7
3.2 Online News Delivery	8
3.3 Online Advertising.....	9
3.4 Online Classified Advertising.....	10
4 News Personalization	11
ARCHITECTURE	14
5 Netpage System Architecture	15
5.1 The Internet	15
5.2 Netpages and Netpage Documents.....	16
5.3 The Netpage Network.....	17
5.4 The Netpage Printer	18
5.5 The Netpage Pen.....	21
5.6 Netpage Interaction	22
5.7 Forms.....	23
5.8 Standard Features of Netpages.....	24
5.9 The Help System	24
6 Personalized Publication Model	26
6.1 Editorial Personalization	26
6.2 Advertising Localization and Targeting	27
6.3 User Profiles	29
6.4 Intelligent Page Layout	30
6.5 Document Format.....	30
6.6 Document Distribution	31
6.7 On-Demand Documents	32
7 Security	34
7.1 Cryptography	34
7.2 Netpage Printer Security.....	35
7.3 Netpage Pen Security.....	36
7.4 Secure Documents	37
7.5 Non-Repudiation.....	37
8 Electronic Commerce Model.....	38
8.1 Secure Electronic Transaction (SET).....	38
8.2 SET Payments.....	38
8.3 Micro-Payments	39
8.4 Transactions	39
APPLICATIONS AND BUSINESS MODELS	40
9 Applications	41


9.1	Personalized Subscriptions	41
9.2	On-Demand Publications	42
9.3	E-Commerce	43
9.4	Communication	44
9.5	Corporate and Government	47
9.6	Personal	48
9.7	Consumer Electronics	50
9.8	Advertising	50
9.9	Product Information and Registration	51
9.10	Geographic Information	53
10	Business Models	54
10.1	System Principles	54
10.2	Bootstrapping	55
10.3	Maturity	56
	REFERENCES	59
11	Silverbrook References	60
12	Other References	61

OVERVIEW

1 Introduction

Netpages are pages of high-quality text, graphics and images printed on ordinary paper, but which work almost like interactive Web pages. Information encoded on each page in invisible ink is picked up by an optically-imaging pen and transmitted to the network. Active “links” and “buttons” on each page can be “pressed” with the pen to request information from the network or signal preferences to a server. Text written by hand on a Netpage is automatically recognized via the pen, allowing forms to be filled in. Signatures recorded on a Netpage are automatically verified, allowing e-commerce transactions to be securely authorized.

The pen, shown on the right, works in conjunction with a Netpage Printer, an Internet-connected printing appliance for home, office or mobile use. The pen is wireless and communicates with the Netpage Printer using an encrypted radio frequency signal.




The Netpage Printer delivers, periodically or on demand, personalized newspapers, magazines, catalogs, brochures and other publications, all printed at high quality as interactive Netpages. Unlike a personal computer, the Netpage Printer, shown on the right, is an appliance typically wall-mounted in the kitchen or near the breakfast table, i.e. the place where the morning news is first consumed, and the household's point of departure for the day. It also comes in tabletop, desktop, portable and miniature versions.

Netpages printed at their point of consumption combine the ease-of-use of paper with the timeliness and interactivity of an interactive medium.

Netpages are crucially enabled by Memjet printing technology [2], which makes high-speed magazine-quality printing affordable to consumers. A Netpage publication has the physical characteristics of a traditional newsmagazine, i.e. a set of letter-size glossy pages printed in full color on both sides, bound together for easy navigation and comfortable handling. The Netpage Printer prints 60 to 90 full-color Netpages per minute.

The Netpage Printer exploits the growing availability of broadband Internet access. Cable service is available to 95% of households in the United States [57], and cable modem service offering broadband Internet access is already available to 20% of these [9]. The Netpage Printer also operates with slower connections, but with longer delivery times and lower image quality.

Netpage Publication Servers on the Internet deliver print-quality publications to Netpage Printers. Periodical publications are delivered automatically to subscribing Netpage Printers via pointcasting and multicasting Internet protocols. Personalized publications are filtered and formatted according to individual user profiles.



A Netpage Printer supports any number of pens, and a pen can work with any number of Netpage Printers. Each Netpage Pen has a unique identifier. A household may have a collection of colored Netpage Pens, one assigned to each member of the family. This allows each user to maintain a distinct profile with respect to a Netpage Publication Server. A Netpage Pen can also be registered with a Netpage Registration Server and linked to one or more payment card accounts. This allows e-commerce payments to be securely autho-

rized using the Netpage Pen. The Netpage Registration Server compares the signature captured by the Netpage Pen with a previously registered signature, allowing it to authenticate the user's identity to the e-commerce system. Other biometrics can also be used to verify identity. A version of the Netpage Pen includes fingerprint scanning, verified in a similar way by the Netpage Registration Server.

Although a Netpage Printer delivers periodicals such as the morning newspaper without user intervention, it never delivers unsolicited junk mail. It only delivers periodicals from subscribed or otherwise authorized sources. The Netpage Printer is unlike a fax machine or e-mail account which is visible to any junk mailer who knows the telephone number or email address.

1.1 PUBLISHING TO NETPAGE PRINTERS

There are a number of advantages to publishing to Netpage Printers. The magazine quality of Netpage Printer output makes it a more attractive publishing and advertising medium than both traditional newsprint and computer screens.

The cost of paper and ink consumption is transferred to the user. Subscription fees can be eliminated entirely in lieu of the user taking on this extra cost, and the user thus perceives a better-value product. The erratic price of newsprint is removed from publishers' profit equations, resulting in more stable margins. A new market for paper and ink consumables, with its own margins, is created.

The cost of consumables can be selectively subsidized, for example when non-editorial publications such as product brochures and account statements are printed.

Capital and maintenance expenditure on printing plant is effectively transferred to the user, although the perceived expense is small because Netpage Printers are sold at close to cost or given away to encourage adoption, subsidized by future advertising profits. Maintenance can also be subsidized or its cost included in a longer term service agreement.

Costly physical distribution is replaced by electronic distribution via a preexisting and widely subscribed network - the Internet.

Both the editorial and advertising content of publications delivered via the Netpage Network can be customized for each user. Editorial content can be personalized according to the user's profile. Advertising can be localized to the user's locality and optionally targeted to the user's demographic.

A personalized publication can be a small fraction of the size of its traditionally-delivered counterpart, yet contain the same amount of information relevant to the user, and in a more accessible form. The user appreciates the more efficient and digestible publication.

Localized advertising can be targeted to more specific localities and their associated demographics, and this allows advertising space to be exploited more efficiently, i.e. with less waste. Advertisers are constantly pressing traditional publishers for greater localization, something which they have great difficulty delivering cost-effectively.

By revealing personal information such as age, gender, marital status, income, profession, education, etc., the user can allow the advertising to be more carefully targeted. In return they can receive greater subsidies and discounted product offers. As advertising becomes more targeted, it becomes less of a nuisance and more of a service in itself.

Although the publication's per-page circulation figures fall drastically, the publication's actual per-section readership is preserved, and the correspondingly higher advertising rates for personalized delivery can exactly compensate for this.

Advertising delivered via the Netpage Network has the dual benefits of print and online delivery. Print supports the impact of large-format ads. Online delivery supports customization, linking, and measurability, and consequently online charging models.

Consider a full-page advertisement for a new car model in a newsmagazine delivered via the Netpage Network. The advertising campaign can be national or even international. The ad only appears if compatible with the user's demographic, either implied by their ZIP Code or more explicitly by their personal details. Anyone who requests a product brochure via the on-ad button receives one immediately via their Netpage Printer, customized with a list of local dealers. If they press a particular dealer's "contact me" button in the brochure, the dealer receives a message via the system and contacts the user by telephone.

The publisher profits in the normal way by selling the advertising space, but can also profit by receiving a fee on the "click-through" to the brochure, and a commission on any product sale which eventuates.

The Netpage Network promises to be the most effective advertising medium ever conceived. It combines the editorial and print quality of traditional publications with arbitrarily finely targeted advertising, and provides a direct link between advertising, product information, and purchasing. Added revenue from click-through fees and e-commerce commissions may even allow users' costs - printer, ink, paper, and Internet access - to be fully subsidized.

2 The Demise of Paper

Online publication has many advantages over traditional paper-based publication. From the consumer's point of view, information is available on demand, information can be navigated via hypertext links, information can be searched, and information can be automatically personalized.

From the publisher's point of view, the costs of printing and physical distribution are eliminated, and the publication becomes more attractive to the advertisers who pay for it because it can be targeted to specific demographics and linked to product sites.

Online publication also has a few disadvantages. Computer screens are inferior to paper. At the same quality as a magazine page, an SVGA computer screen displays only about a fifth as much information¹. Both CRTs and LCDs have brightness and contrast problems, particularly when ambient light is strong. Ink on paper, being reflective rather than emissive, is both bright and sharp in ambient light.

Faced with reading more than the most trivial amounts of text on a screen, most people prefer to print it before reading it. Increasingly, online publishers are recognizing this and providing information in formats suitable for printing. At one extreme this means providing text-only versions of documents so they print efficiently, i.e. without imposing a screen format on the printed page; at the other extreme it means providing formatted versions of documents - e.g. in Adobe's Portable Document Format (PDF) - so they print at high quality.

Editorial content is often compromised to fit the online medium and the habits of the people who frequent it, who tend to browse rather than read. Although powerful new advertising models become possible, it becomes more difficult to deliver effective advertisements.

To truly enable online publication, many people envisage a universal information appliance - a lightweight portable "tablet" with a page-size touch-sensitive color display and a high-bandwidth wireless connection to the Internet. First proposed by Xerox PARC's Alan Kay in the mid 1970s in the form of the "Dynabook", and partially realized in recent paperback-sized electronic books, even Bill Gates is now confidently predicting that such a device will soon augur the death of paper publications [25].

To achieve low power consumption, low weight, and paper-like display quality, a bistable reflective display technology is required. Several candidates are now emerging from the labs, including Kent Display's cholesteric LCD technology (chLCD) [36], Xerox' "Gyricon" rotating ball technology [28], and E Ink's electrophoretic technology [21,47]. ChLCD is arguably closest to practical deployment [20].

Next-generation cellular phone networks promise 2Mbps packet switching [22], comparable to the broadband access people are getting used to in cable networks [44]. Satellite networks, while offering or promising still higher speeds [29,50,63], require receivers difficult to deploy in mobile devices.

Beyond the vision of the basic tablet, E Ink imagines its digital ink "printed" onto a number of flexible pages bound into a book, preserving the physical navigability of a paper-based publication, and approaching its low cost, but allowing the pages to be rewrit-

1. Assuming a magazine page has an equivalent digital resolution of 200 continuous-tone pixels per inch.

ten electronically in place. They optimistically predict newspapers delivered in this way within five years [16], despite fundamental problems yet to be overcome [20].

The advantages of a tablet are many. Unlike a desktop or notebook computer, a tablet may actually provide a pleasant reading experience. Unlike a paper publication, a tablet provides intelligent access to an unlimited amount of information; its weight is not dictated by the amount of information it carries. More than just an information appliance, it can also act as a multi-purpose multimedia communications device and interactive entertainment device.

A tablet has disadvantages too. It uses batteries which run down and have to be recharged. It may break when dropped or malfunction when exposed to hot coffee. It's not quite cheap enough to be disposable - so there's still a problem if it's misplaced or stolen. It has a "user interface" which has to be learned. The leading candidate display technology - chLCD - is still less than half as reflective (i.e. "bright") as paper.

The drawbacks of traditional paper-based publications have little to do with paper itself, and much to do with how the information gets onto the printed page. The economics of centralized printing and distribution prevent the kind of information selection, personalization and navigation people have come to expect from interactive electronic media such as the Internet. The inefficiency of printing and distributing a hundred-page newspaper to a customer who may read only a few pages is widely decried.

Given a technology such as Memjet, it becomes economic to print high-quality publications at their point of consumption rather than at their point of production. The Netpage Printer leverages Memjet to deliver personalized publications to the home, gaining many of the advantages of online publication, while retaining the ease-of-use of high-quality printed ink on paper.

Netpages and the Netpage Printer address the key problems of online publication, without relying on the development and consumer acceptance of a new reading device.

3 News and Advertising Trends

3.1 NEWSPAPERS

People obtain news from a variety of sources - network and cable television, radio, daily newspapers, and weekly newsmagazines. In the United States, although the various news media are healthy and profitable, per capita news consumption is somewhat in decline as a new generation of young adults have less time to read and favor television entertainment over news [17]. Yet six out of ten adult Americans read a newspaper every day [53].

The United States has about 1500 daily newspapers with a total circulation of 57 million. Just the top ten "national" dailies (Wall Street Journal, USA Today, New York Times, Los Angeles Times, Washington Post, etc.) account for a circulation of 10 million. The major weekly newsmagazines (Time, Newsweek, U.S. News) have a similar (weekly) circulation of about 10 million.

In 1997, newspaper companies' revenue exceeded \$24 billion, a five-year high, and margins nudged 20% [39], due both to increased spending on advertising and to reduced prices of newsprint.

Television and radio, by their nature, excel at delivering breaking news. Newspapers and newsmagazines, on the other hand, deliver the depth and analysis behind the headlines. Broadcast news in isolation does a poor job of informing the public. The more local the news is, the poorer the broadcast coverage, and the greater the public's dependence on newspapers.

Newspaper content and packaging has evolved considerably since the 1970s. News is somewhat softer, news stories are shorter and more well-written, there are more feature articles, and there is more editorial and reader opinion. Newspapers are more structured. Identifiable sections make them more accessible, and provide greater focus for advertisers. Much special-interest content has migrated from daily inclusion to weekly sections. These cover topics such as lifestyle, personal finance, entertainment, technology, etc. The proportion of graphics and pictures is greater. Color is widely used. Newspapers are easier to use and more entertaining than ever before, if at the expense of some "hard" news.

Daily newspapers are growing increasingly dependent on the various wire services. A newspaper may excel at local and regional news, but rely on the major wires (Associated Press and United Press International) for national and international news, the so-called "supplemental" wires (LA Times/Washington Post, NY Times, Scripps-Howard, Knight-Ridder Tribune, etc.) for specific strengths (and value-for-money), and the international wires (Reuters etc.) for international perspective. A growing number of newspapers operate more as news aggregators than news gatherers.

Advertising typically contributes more than 75% of newspaper and magazine revenue, while subscriptions contribute less than 25% [27,51,52]. National advertising makes up roughly 14% of advertising spending, retail advertising 46%, and classified advertising 40% [52].

Advertisers are pursuing increasingly specific targeting, favoring quality newspaper readership over raw circulation [24], and using more targeted media where possible. Magazines, for example, have more specific readerships than newspapers, free "shoppers" are

very localized by their nature, while direct mailers can target demographics based on individually-categorized ZIP Codes, or databases of individuals.

Newspapers have responded with geographically zoned editions to support local advertising, and greater sectioning of their product. They have also expanded their page counts to provide more advertising scope, despite erratic newsprint prices in the 1990s [5].

Despite this, there is ongoing conflict between newspapers' mass distribution model, and advertisers' need for micro-targeting [26]. This conflict, coupled with advertisers' desire for higher-quality printing of color images, is motivating a shift from run-of-press (ROP) advertising to inserts [52]. The downside to inserts is that editorial context is lost.

3.2 ONLINE NEWS DELIVERY

Fearing the online migration of advertising, traditional news publishers from both broadcast and print have ventured into Internet-based news delivery, wanting to establish a presence at whatever cost before newcomers become entrenched. Most newspapers are still reporting losses from their online operations [52].

Online news delivery offers a number of advantages. Breaking news can be delivered as soon as it happens. News can be customized for individual readers according to their preferences and geographic locations. Readers can explore stories to arbitrary depth, follow links to related resources, and search archival material. Readers can participate in discussion groups and contribute to opinion polls. The news itself can incorporate audio and video clips, and can include live transmissions, converging with broadcast.

Online news delivery also has disadvantages. Computer screens are of limited size and quality compared with print. Few people enjoy reading a story of any length on a computer screen. Computers are not portable in the wide sense that a newspaper is. The news may be more timely, but the time and place in which it can be consumed are more constrained than with a newspaper.

Despite the power of hypertext, many online readers express a preference for a linear presentation, "where they [can] skim one section after another until the presentation [is] exhausted" [11]. Interestingly, a majority of traditional newspaper readers admit they scan every page in the main section of the newspaper [53], looking for items of interest without necessarily knowing what they're looking for, and achieving some kind of closure at the end. Online hypertext, by contrast, is both a limitless resource and a bottomless pit.

While traditional news publishers such as The New York Times can deploy full editorial content online [64], newcomers such as Yahoo typically only provide "raw" news items sourced from the wire services [48].

A recent survey indicates that 21% of the 74 million Internet users in the United States regularly read news online as an alternative to traditional print and broadcast sources, and 16% obtain a major proportion of their news online [54]. More broadly, between 37% and 64% of the Internet population reads news online at least once a week. The fluctuations in the figures are related to what may be happening in the news. Major or breaking news stories attract more users - 46% of Americans say they only follow national news stories when "something major is happening" [23].

With 41% of Americans online, the Internet population has become mainstream, and the weather has become the most popular news online. This is closely followed by technology

news, entertainment news, and local news. As one observer puts it, all of this “sound[s] like the 6 o’clock news” [54]. As a reflection of these habits, the online audience share of national newspapers has diminished from 23% in 1995 to 16% in 1998, while the online audience share of broadcast TV sites has grown.

3.3 ONLINE ADVERTISING

At its simplest, advertising alerts a motivated customer to the availability of a product, possibly at a competitive price. At a more sophisticated level, advertising seeks to influence future purchasing decisions by creating brand awareness. Ultimately, advertising seeks to create desire for a product even when actual need is absent.

Advertising prices are traditionally based on how many people see the advertisement, and their spending power in relation to the product. In practice, the more homogeneous the demographics of the audience, the easier it is to match to a product, and hence the higher the corresponding advertising cost per thousand (CPM). Broadcast media use ratings and timeslot demographics to set advertising rates. Print media use audited circulation figures and sectional readership demographics.

The simplest online advertising model is also based on how many people see the ad. Online this has the advantage of being based on solid numbers, since the number of “impressions” of a particular Web page can be counted exactly.

The specific advantage of an online ad, however, is that the ad itself can measurably capture a sales lead by acting as a link to a product site. The product site may simply provide more product information in the form of specifications, pricing, and ordering details. It may also support immediate online ordering, thus completing the link from ad to sale. Beyond providing simple ad exposure, it is this measurable linking of advertisement to sales lead or sale which is the strength of online advertising [55]. Cost per click (CLC) charging is gaining acceptance but is still controversial.

Beyond CLC, there exists the possibility of paying a commission to the ad host on any sale that actually eventuates [58]. Amazon.com is probably the best-known example of a company paying commissions to other sites in this way.

The broader advantage of online advertising is that advertising can be localized and targeted arbitrarily finely, in conjunction with the publication of online content such as news. This is the strategy pursued by online advertising agencies such as Click-Through [15], which acts as the middle-man between advertisers and online content publishers. They expect online advertising to represent more than 10 percent of all advertising revenue by 2001.

Since online ads are necessarily small-format, they communicate best with motivated customers already on the look-out for a particular product or service. Online ads are less suited to building brand awareness or creating buying desires, since the real substance of the advertisement - the product Web site - is a click away from the initial small-format ad. A small-format online ad can’t provide the single-hit emotional impact of a large-format print ad, and conversely, the online world can’t support the large-format ads that print can.

So-called interstitial ads, which appear full-screen when traversing from one page of information to the next, go some way to providing a medium for larger-format ads online [58]. User resistance, however, seems to be preventing their widespread use.

3.4 ONLINE CLASSIFIED ADVERTISING

Classified advertising is indisputably suited to online delivery. Unlike their traditional printed counterparts, online classifieds can be easily searched, and are not subject to space constraints. The online migration of classified advertising is considered a serious threat to newspapers' classified advertising profits [60], and some newspapers are building an online presence for this reason alone. Some observers predict as much as 50% of classified advertising revenue moving online within the next ten years [72].

Another problem faced by newspapers, who rely on classifieds for up to 40% of advertising spending, is that many newcomers are offering free online classified advertising as a way of building a venue for non-classified advertising.

4 News Personalization

From the reader's point of view, a personalized news publication can provide more information in fewer pages. The actual form this personalization takes, however, is not necessarily obvious.

The MIT Media Lab's News in the Future (NiF) project has been championing the concept of "The Daily Me" for almost two decades [45]. Nicholas Negroponte, one of the project's founders, envisages a highly personalized news publication which is no longer driven by "what other people think is news" [49]. By way of examples close to his own needs, it includes news about people and places about to be encountered, and puts "the most important [news] of all" - a summary of e-mail - on the front page [7]. Negroponte recognizes the need to vary the degree of personalization, advocating a higher "serendipity factor" on a lazy Sunday than on a working weekday.

The opposing view holds that the value of a news publication lies precisely in its *shared* nature. It reflects the common concerns and values of a community of readers, and establishes a baseline of expectations of what they are all supposed to know [67]. As a consequence, the publication also speaks with a consistent editorial voice and with consistent assumptions about the reader's level of background knowledge. Such a shared publication allows its readers to orient themselves in relation to their community.

NiF's Walter Bender answers the charge (in his own words) of "The Daily Me engendering a fragmented world populated by self-interested myopes", by stressing the possibilities of personalizing individual news items [6]. This can consist in varying the depth of an item, or supplementing it with background information, based on the reader's level of knowledge. It can involve interpreting information relative to the reader's background, such as (somewhat dubiously) making value judgements about the weather relative to the reader's normal home town weather. It can also be as simple as using metric rather than imperial units.

FishWrap [11,46], MIT's personalized campus newspaper and NiF's latest offering, goes further by creating a front page whose content represents an explicit community consensus. Each front page news item is prioritized according to the number of readers who put it forward for inclusion. The rest of the newspaper is still personalized according to each reader's profile, consisting of reader-defined sections containing topics of interest.

There are two implications of recognizing the shared nature of news. Firstly, some news is news to everybody in a community, no matter how personalized they claim they would like their news to be. This implies that the community must make decisions about news item priority, either directly (as in FishWrap) or indirectly via a proxy (i.e. an editor).

Secondly, a news item can only be properly understood in the context of the community for which it is intended. This implies that a news item must be branded with its source (assuming that the source implies the intended target). As an example, it is significant whether a news item regarding the proof of Fermat's last theorem is branded with New Scientist or The New York Times. To a professional mathematician, the latter implies, by its very existence, that the proof is of significance beyond the scientific community.

Of course, a news item must also be branded to allow its source to build and maintain its brand. The brand then allows the reader to infer the quality of the news item from the known quality of the source.

Most personalization of news uses *feature-based filtering*. This means that news item content is matched to topics and keywords in the reader's profile. News sources tag the items they produce with various information to allow them to be effectively filtered. This tagging may be brief or extensive, and may include such things as news item urgency, byline, news category, subject(s), keyword(s), date and time, and location [33]. The body text of a news item can also be scanned directly for keywords, but this may result in false matches if keywords are interpreted out of context. Items in the text such as personal names and locations can be tagged to reduce such ambiguity [33]. Similarly, dates, times, and monetary amounts can be tagged to allow localized presentation.

Feature-based filtering suffers from a number of problems. Filtering based on tags is only as good as the original tagging. The latest tagging standards are only just beginning to be adopted [8]. Filtering based on the text itself is constrained by the intelligence of the text parsing. If based simply on keyword matching, it can be both inaccurate, generating false matches because of word sense ambiguity, and imprecise, generating false mismatches because of a lack of inference.

Feature-based filtering is incapable of discerning more abstract attributes such as quality, style, and point-of-view (unless they're indicated by tags). And since it only matches items anticipated by the user's profile, it is a poor generator of serendipitous finds.

FishWrap's front page comes into existence based on a crude form of *collaborative filtering*. In its broader form, collaborative filtering involves sharing recommendations (or ratings) among *like-minded* people [56]. This means that one person's ratings influence another person if and only if the two share similar interests, i.e. they have similar rating histories. Collaborative filtering overcomes many of the problems of feature-based filtering, since ratings originate with people who have digested the items in question, rather than from automated analysis of the items. Collaborative filtering sidesteps the issue of *explaining* why a person might like a particular item.

Collaborative filtering has problems of its own. The system only works if people are willing to contribute ratings. In contributing ratings, of course, they are both doing the community a service and tuning their own interest profiles. The statistical error in correlating people's interests decreases as the number of ratings increases. However, incentives may have to be offered to encourage people to contribute ratings.

To bootstrap the accumulation of ratings for new items, an independent mechanism must exist to distribute them to a critical mass of people. Conversely, to bootstrap the accumulation of interest profiles for new users, an independent mechanism must exist to distribute a critical mass of items to them.

To allow meaningful accumulation of ratings, a sufficient period of time must be allowed to elapse. This may conflict with the timely delivery of items in question.

The statistical correlation between different people's interests, represented by their rating histories, is most meaningful when the ratings apply to homogeneous items. For a set of heterogeneous items, collaborative filtering is best applied to homogeneous subsets.

In a news setting, collaborative filtering is best applied to feature articles. Features have the longer life span required to support the accumulation of ratings, and are often appreciated for abstract qualities best singled out by collaborative filtering (good writing, humor, incisiveness, etc.).

Naturally, the larger a publication's readership, and the better its taste in relation to a potential reader, the stronger the publication's brand will appear to that reader.

Although it's easy to become preoccupied with automatic filtering, in reality there's more to editing the news than just filtering news feeds. An editor also solicits news, commissions analysis, and offers opinion, ideally ensuring that the publication offers a balanced and complete view of the world.

Perhaps the most important personalization step a reader takes is in selecting a particular publication from a set of available publications, based on its perceived quality and relevance.

Thus the publication's brand equates to the highest-level and most useful filter of all.

ARCHITECTURE

5 Netpage System Architecture

5.1 THE INTERNET

The Internet is a worldwide collection of interconnected networks which communicate using the TCP/IP protocol suite [59]. A TCP/IP-based internetwork not connected to *the* Internet is often referred to as *an* internet (i.e. with a lower-case 'i'). When an internet is deployed within an organization, it is often termed an intranet.

Access to the Internet is widespread in developed countries. In the United States, for example, 41% of the population has access to the Internet [54].

While most consumers still access the Internet via low-speed dial-up modems connected to the switched telephone system, inexpensive broadband access is becoming available to a growing number of households via the cable networks. Cable service is available to 95% of American households [57], and cable modem service is available to a 20% subset [9]. While dial-up modems offer speeds of up to 56Kbps, cable modems offer practical speeds of up to about 3Mbps¹, i.e. over 50 times faster.

DSL (Digital Subscriber Line) [12,13], while offering similar speeds to cable modems but via the telephone system, is not yet widely used. ISDN (Integrated Services Digital Network), although widely used for corporate access, has had little consumer impact due to its high price and comparatively low performance.

The deployment of third-generation (3G) cellular telephony within the next few years will bring practical mobile broadband speeds of 2Mbps [22]. 3G cellular uses WCDMA (wide-band code-division multiple access), a spread-spectrum technology. Satellite systems are arguably closer to offering even faster broadband Internet access [29,50,63].

The core of the Internet is made up of a number of independent high-speed fiber-optic networks connected into NAPs (Network Access Points) or peered directly. These have until recently used single-wavelength TDM (Time-Division Multiplexing) SONET (Synchronous Optical Network) transmission systems which utilize about 1% of an optic fiber's capacity to yield a 2.5Gbps OC-48 channel². Carriers are now beginning to deploy multi-wavelength DWDM (Dense Wavelength-Division Multiplexing) systems which yield up to 40 such channels per optic fiber, thus increasing network capacity significantly without requiring the laying of more fiber [10,40]. Internet architects are therefore now contemplating aggregate capacity in the terabit (Tbps) range.

The Internet uses the four-layer TCP/IP protocol suite. The application layer provides various end-to-end application services, and is a client of the transport layer which provides end-to-end delivery services. The transport layer in turn is a client of the network layer which provides packet routing. The network layer is a client of the link layer which encapsulates specifics of the protocols and hardware of the actual communications links.

The core Internet transport protocol, TCP (Transmission Control Protocol), provides a reliable end-to-end delivery service. The core Internet network protocol, IP (Internet Protocol), provides an unreliable and connectionless packet routing service. IP may lose or

1. Although the cable supports 30Mbps and the cable modem theoretically supports 10Mbps.

2. SONET channels have an OC-*n* designation, where OC stands for Optical Carrier, and *n* gives the channel speed in units of about 52Mbps. An OC-48 channel therefore has a speed of about 2.5Gbps.

deliberately discard packets, and may deliver packets out of order, and it is the responsibility of a higher layer to provide a reliable end-to-end service.

With the proliferation of streaming media services on the Internet, support for multicast is spreading rapidly. Multicast is a form of broadcast with a specific set of recipients. It makes efficient use of network capacity because a packet traverses a network link once rather than once per recipient. It is particularly efficient if the recipients are connected to the Internet via an intrinsically broadcast medium such as cable or satellite. The @Home cable network has successfully enabled multicasting of streaming media services [38].

IP Multicast is an extension of IP, and so is unreliable. While this is often acceptable for time-critical data such as streaming video, it may not be acceptable for other shared data types. Significant effort is being expended to develop reliable multicast transport protocols on top of IP Multicast. Although several reliable multicast protocols are available and have been deployed [41,42,30,31], the Internet standardization process is incomplete [32].

5.2 NETPAGES AND NETPAGE DOCUMENTS

Netpages are the foundation on which a Netpage Network is built. They provide a paper-based user interface to published information and interactive services.

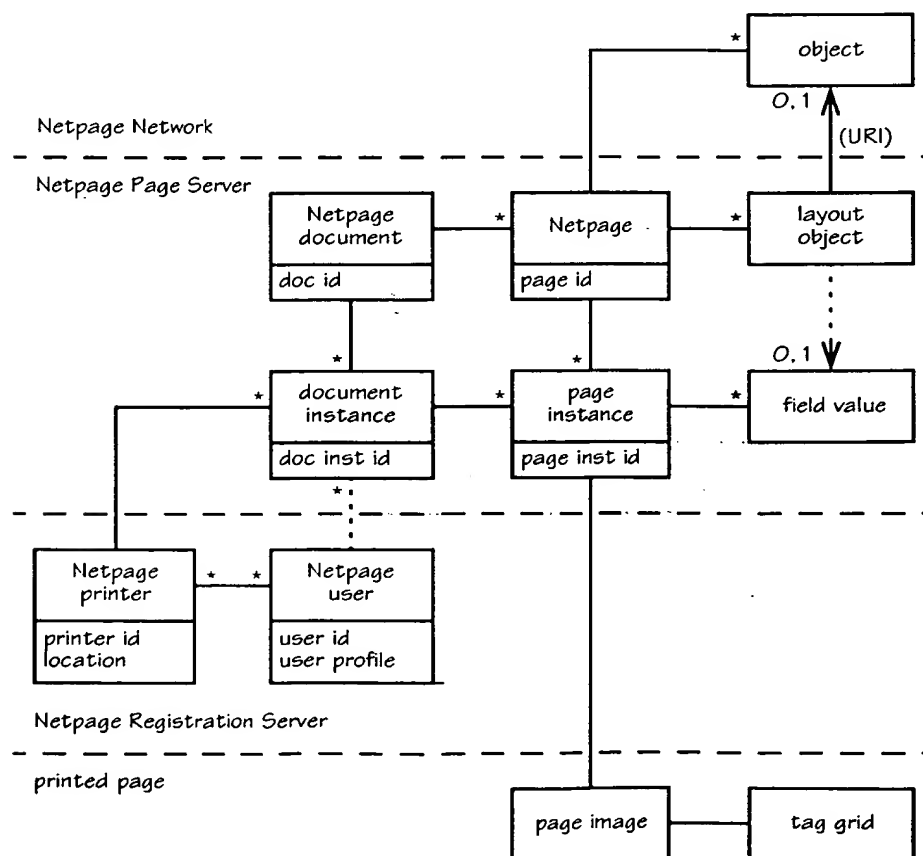


Figure 1. Netpage Document structure (* indicates an n-ary relationship)

Each Netpage consists of a compact page layout maintained persistently by a Netpage Page Server. The page layout refers to objects such as images, fonts and pieces of text, typically stored elsewhere on the Netpage Network.

Netpages are organized into Netpage Documents. Both Netpages and Netpage Documents are assigned globally unique identifiers.

Each Netpage Document has a set of document instances, each of which describes a printed instance of the document. Each Netpage in the Netpage Document has a corresponding set of page instances, each of which describes a printed instance of the page. Both page instances and document instances are assigned globally unique identifiers. They are also uniquely associated with the printer on which they are printed and the user who initiated the print request, if known.

Each page instance maintains a set of user-supplied values for fields in the page layout. This ensures that user input is captured and stored independently for each page instance. The separation of page instances and Netpages is crucial for pages which contain input fields, i.e. forms. It is not crucial for pages devoid of input fields, but still useful because it supports independent auditing of each page instance.

The physical page image includes encoded information which identifies the page instance and hence the Netpage to which it corresponds. It also includes encoded information which superimposes an addressable spatial grid over the page image, to allow pen actions performed relative to the page image to be correlated with the contents of the page layout.

The encoded information is normally printed in infrared-absorptive ink on any normal paper substrate which is infrared-reflective. Near-infrared wavelengths are invisible to the human eye but are easily sensed by a solid-state image sensor with an appropriate filter.

The encoded information is picked up by an infrared-imaging pen and transmitted to the associated Netpage Printer. The pen is wireless and communicates with the Netpage Printer using an encrypted radio frequency signal.

The encoded information is organized as a set of tags, each containing both the id of the page instance and the position of the tag. The tags tile the entire page image, and are sufficiently small and densely arranged that the pen can reliably image at least one tag even on a single click on the page. It is important that the pen recognize the page instance id and position on every interaction with the page, since the interaction is stateless.

The tags are error-correctably encoded to make them resilient to errors introduced by dirt on the page or during the imaging process.

Memjet-based Netpage Printers are designed to print a tag grid using infrared (IR) ink. Printers not enabled for IR printing have the option to print tags using IR-absorptive black ink, although this restricts tags to otherwise empty areas of the page. Although such pages have more limited functionality than IR-printed pages, they are still classed as Netpages.

5.3 THE NETPAGE NETWORK

A Netpage Network consists of a distributed set of Netpage Publication Servers, Netpage Page Servers, and Netpage Printers connected via an internet. In technological terms this document describes *any* Netpage Network. In business terms it usually refers to *the* Netpage Network connected via *the* Internet.

As described above, a Netpage Page Server maintains persistent information about Netpage Documents, Netpages, and their printed instances, to allow pen operations on printed pages to be interpreted intelligently.

The Netpage Network includes any number of Netpage Page Servers, each handling a subset of Netpages. As described above, each page instance is identified by a globally unique id which is encoded in the tag grid of the corresponding printed page. The Netpage Printer uses this id to retrieve the page layout of the page from a Netpage Page Server when it needs to interpret pen operations relative to the page.

The Netpage Printer uses the internet Distributed Name System (DNS) to resolve a Netpage instance id into a page instance maintained by a particular Netpage Page Server.

The DNS is a protocol and a hierarchical system of name servers used to resolve internet domain names into resources. Planned enhancements to the DNS allow it to be used to resolve more general Uniform Resource Identifiers (URIs), and in particular Uniform Resource Names (URNs), into resource locations [19]. Netpage instance ids are formulated as URNs, allowing the enhanced DNS to be used to resolve them. In the absence of timely standardization and deployment of an enhanced DNS on the Internet, the Netpage Network can deploy its own system of enhanced name servers.

A Netpage Publication Server is an internet server which publishes Netpage Documents to Netpage Printers. It is described in Section 6.

5.4 THE NETPAGE PRINTER

The Netpage Printer is the appliance which prints Netpage Documents. It is connected to a Netpage Network via an internet, ideally via a broadband connection.

Apart from identity and security settings in non-volatile memory, the Netpage Printer contains no persistent storage. As far as a user is concerned, *the network is the computer* [61]. Netpages function interactively across space and time with the help of the distributed Netpage Page Servers, independently of particular Netpage Printers.

The Netpage Printer receives Netpage Documents from Netpage Publication Servers. Each document is distributed in two parts: the page layouts, and the actual text and image objects which populate the pages. Because of personalization, page layouts are typically specific to a particular subscriber and so are pointcast to the subscriber's printer. Text and image objects, on the other hand, are typically shared with other subscribers, and so are multicast to all subscribers' printers.

The Netpage Publication Server optimizes the segmentation of document content into pointcasts and multicasts. After receiving the pointcast of a document's page layouts, the printer knows which multicasts, if any, to listen to.

Once the printer has received the entire document's page descriptions, i.e. page layouts and objects, it can print the document.

The printer rasterizes and prints odd and even pages simultaneously on both sides of the sheet. It therefore contains duplexed print engines and imaging units.

The printing process consists of two decoupled stages: rasterization of page descriptions, and expansion and printing of page images. The raster image processor (RIP) consists of

one or more standard DSPs running in parallel. The duplexed print engines consist of custom processors which expand, dither and print page images in real time, synchronized with the operation of the printheads in the imaging units.

There are four major design variations embodied in the various Netpage Printer models:

- *form factor*: pocket, portable, desktop, wall-mount or tabletop
- *printhead width*: 4" (photo), 8½" (portrait Letter) or 11" (landscape Letter)
- *paper source*: cut sheet or print cartridge
- *Internet connection*: wired or wireless

The form factor variations yield five basic models, each with variants determined by printhead width (and hence printing speed), and paper source. Eight planned models are defined in Table 1, and illustrated in Figure 2.

Table 1. Netpage Printer models

model	form factor	variant	printhead width	paper source	Internet connect
Microprinter	pocket	R	4"	cartridge	wireless
Travelprinter	portable	R	8½"	cartridge	wireless
Deskprinter	desktop	R	11"	cartridge	wired
Wallprinter	wall-mount	-		cut sheet	or
		Pro		cut sheet	wireless
		Pro R		cartridge	
Tableprinter	tabletop	Pro		cut sheet	
		Pro R		cartridge	

The Deskprinter, Wallprinter and Tableprinter models can be factory-configured with various network modules, allowing both wired and wireless versions. The Microprinter and Travelprinter both use a cellular telephone module, with the promise of broadband speed within a few years.

The Wallprinter models are ideal for unobtrusive installation in a home, while the Tableprinter models might be preferred in an office environment. Note that the Tableprinter models are Wallprinter models factory-adapted for tabletop use via a stand. The Deskprinter, with its small footprint, is ideal for both home and office use.

The Microprinter prints normal Netpages at quarter size, and provides full wireless Netpage Network access in a pocket device.

The paper roll cartridge contains both paper and ink. The paper is in the form of a continuous roll, cut on demand by the printer. The 11" print cartridge has a capacity of 1000 Letter sheets. It also contains the glue supply for binding the sheets of a document together. The 8½" print cartridge has a capacity of 50 Letter sheets, or equivalently 100 A5 sheets. The 4" print cartridge has a capacity of 36 6x4 photos, or 41 quarter-size Netpages. The 8½" and 4" print cartridges don't contain a glue supply because neither the Microprinter nor the Travelprinter includes a binding mechanism.

The 4" printhead models print at 30 quarter-size pages per minute. The 8½" printhead models print at 60 pages per minute, or 30 duplex sheets per minute. The 11" printhead models print at 90 pages per minute, or 45 duplex sheets per minute.




Figure 2. Netpage Printer family

5.5 THE NETPAGE PEN

The Netpage Pen operates both as a normal marking ink pen and as a non-marking stylus. When either nib is in contact with a Netpage, the pen continuously monitors its movements relative to the page. The nib is attached to a pressure sensor. The pen pressure can be interpreted relative to a threshold to indicate whether the pen is “up” or “down”. It can also be interpreted as a continuous value, for example when the pen is capturing a signature, to allow the full dynamics of the signature to be verified.

The pen determines the position of its nib on the Netpage by imaging, in the infrared spectrum, an area of the page in the vicinity of the nib. It decodes the nearest page id and position tag, and adjusts the position given by the tag to account for the distance between the area imaged and the actual nib, and the position of the tag in the imaged area. Although the position resolution of the tag may be low, because the tag density on the page is inversely proportional to the tag size, the adjusted position resolution is quite high, and easily exceeds the minimum 200 dpi resolution required for handwriting recognition [62].

Pen actions relative to a Netpage consist of a series of strokes. A stroke consists of a sequence of time-stamped pen positions on the page, initiated by a pen-down event and completed by the subsequent pen-up event. A stroke is also tagged with the page id of the Netpage whenever the page id changes, i.e. just at the start of the stroke under normal circumstances.

The position tags on the Netpage contain various control bits. One of these instructs the pen to activate its “active area” LED. Thus a region on the page which corresponds to the active area of a button or hyperlink can be encoded to activate this LED, giving the user visual feedback that the button or hyperlink is active when the pen passes over it. Another control bit instructs the pen to capture continuous pen pressure readings and tag the stroke with these readings. Thus a region on the page which corresponds to a signature input area can be encoded to capture continuous pen pressure.

Whenever the pen is within range of a printer with which it can communicate, the pen slowly flashes its “online” LED. When the pen fails to decode a stroke relative to the page, it momentarily activates its “error” LED. When the pen succeeds in decoding a stroke relative to the page, it momentarily activates its “ok” LED.

The pen also contains a pair of passive accelerometers mounted at right angles to each other in the plane normal to the pen’s axis. The accelerometers respond to gravity and allow the pen to compute its tilt. This in turn helps it auto-focus its optics and compute the nib-to-tag displacement. If the stroke is being tagged with pen pressure readings, then it is also tagged with tilt readings.

A sequence of captured strokes, whether tagged with pen pressure and tilt or not, is referred to as *digital ink*. Digital ink forms the basis for the digital exchange of drawings and handwriting, for on-line recognition of handwriting [62], and for on-line verification of signatures.

The pen is wireless and transmits digital ink to the Netpage Printer using a radio frequency signal. The digital ink data is encrypted for security and packetized for efficient transmission, but is always flushed on a pen-up event to ensure timely handling in the printer.

When the pen is out-of-range of a printer it buffers digital ink in internal memory, which has a capacity of more than 12 minutes of continuous handwriting. When the pen is once again within range of a printer, it transfers any buffered digital ink.

A pen can be registered with any number of printers, but because all state data resides in Netpages both on paper and on the network, it is largely immaterial which printer a pen is communicating with at any particular time.

5.6 NETPAGE INTERACTION

When the Netpage Printer receives a digital ink stroke from the pen, it retrieves the page layout of the Netpage identified in the stroke, to allow it to correctly interpret the stroke. The printer resolves, via the DNS, the address of the Netpage Page Server which holds the page layout, and then retrieves the page layout from the server. If the page was recently identified in an earlier stroke, then the printer may already have the address of the relevant Netpage Page Server in its cache. It may also have the page layout itself in its cache, in which case there may be no need to retrieve it.

Once the printer has the page layout of the Netpage to which the pen stroke refers, it can interpret the stroke in relation to the layout and content of the page. This involves hit-testing the objects on the page to determine which objects the pen is interacting with, in much the same way that mouse movements and button presses are interpreted in a graphical user interface system.

A “click” is a stroke where the distance between the pen down position and the subsequent pen up position is less than some small maximum. An object which is activated by a click requires a click to be activated, i.e. a longer stroke is ignored. The failure of a pen action, such as a “sloppy” click, to register is indicated by the lack of response from the pen’s “ok” LED.

There are two kinds of interactive objects on a Netpage: hyperlinks and form fields.

When a hyperlink is activated, the printer sends a request to a handler somewhere on the network. The handler is identified by a URI, and the URI is resolved in the normal way via the DNS. There are three types of hyperlinks: general hyperlinks, form hyperlinks, and selection hyperlinks. A general hyperlink may implement a request for a linked document, or may simply signal a preference to a server. A form hyperlink submits the corresponding form to a form handler. A selection hyperlink submits the current selection to a selection handler. If the current selection contains a single-word piece of text, for example, the selection handler may return a single-page document giving the word’s meaning within the context in which it appears, or a translation into a different language. Each hyperlink type is characterized by what information is submitted to the handler.

Form fields come in four varieties: checkboxes, text areas, digital ink areas, and signature areas. A checkbox accepts a true or false value. Any mark (a tick, a cross, a stroke, a fill zigzag, etc.) captured in a checkbox area is assigned as a true value to the corresponding field. A text area accepts a text string. Any digital ink captured in a text area is automatically converted to text via on-line handwriting recognition and the text is assigned to the corresponding field. A digital ink area accepts raw digital ink. Any digital ink captured in a digital ink area is assigned to the corresponding field. A signature area accepts a hand-written signature. Any digital ink captured in a signature area is automatically verified and the resulting signature token is assigned to the corresponding signature field. Signature verification is discussed in more detail in Section 8.

“Editing” commands, such as strike-throughs indicating deletion, are also recognized in form fields.

Table 2. Summary of pen interactions with a Netpage

object	type	pen input	action
hyperlink	general	click	submit action to handler via URI
	form	click	submit form to handler via URI
	selection	click	submit selection to handler via URI
form field	checkbox	any mark	set field value to true
	text area	handwriting	convert digital ink to text; assign text to field
	digital ink area	digital ink	assign digital ink to field
	signature area	signature	verify digital ink signature; assign signature token to field
none	-	circumscription	convert digital ink to region; select object(s) in region

Because the handwriting recognition algorithm works “on-line” (i.e. with access to the dynamics of the pen movement), rather than “off-line” (i.e. with access only to a bitmap of pen markings), it can recognize run-on discretely-written characters [62] with high accuracy, without a writer-dependent training phase.

Digital ink, as already stated, consists of a sequence of strokes. Any stroke which starts in a particular object’s active area is appended to that area’s digital ink stream, ready for eventual interpretation. Any stroke not appended to an object’s digital ink stream is appended to the remaining inactive area’s digital ink stream.

Digital ink captured in the inactive area is interpreted as a selection gesture. Any circumscription of one or more objects is interpreted as a selection of the circumscribed objects.

The printer maintains a current selection for each pen. The selection contains the most recent object selected, resolved with reference to the page layout and content. The selection can be attached to or pasted into another form, or in general be submitted to a selection handler as described earlier. The selection is cleared after an inactivity time-out to ensure predictable behavior.

Table 2 provides a summary of pen interactions with a Netpage.

5.7 FORMS

As described in Section 5.2, user input on a physical Netpage is ultimately recorded persistently by a Netpage Page Server together with the corresponding page instance. To ensure efficient capture of user input, the printer accumulates input locally. To prevent update anomalies, however, the printer temporarily obtains exclusive access to the page instance from the Netpage Page Server. The printer flushes input back to the server and relinquishes exclusive access when the user initiates a non-local action on the page; after an inactivity time-out on the page; when the printer wishes to free up local storage consumed by the page; and on request from the server.

When the printer submits a form to a form handler, it simply submits the document instance of the form. The form handler retrieves the field values from the Netpage Page Server at its leisure.

A form can also act as a shared “blackboard” between the user and the form handler, i.e. the form handler can query the contents of the form fields maintained by the Netpage Page Server without the user explicitly submitting the form.

For text areas, the raw digital ink is optionally also stored with the page instance on the Netpage Page Server. This allows the form handler to interrogate the raw digital ink should it suspect the original recognition of the handwriting. This might involve human intervention at the application level for forms which fail certain application-specific consistency checks. As an extension to this, the entire background area of a form can be designated as a digital ink area. The form handler can then decide, on the basis of the presence of digital ink outside the explicit fields of the form, to route the form to a human operator, on the assumption that the user may have indicated amendments to the filled-in fields outside of those fields.

Form fields can optionally be tagged to indicate their meaning. Fields tagged in this way may include name and address fields, for example. This semantic tagging allows these fields to be automatically filled in whenever a “blank” form is requested by an identifiable user, i.e. a user who has registered their identity with the system and linked it to the identity of their pen.

5.8 STANDARD FEATURES OF NETPAGES

Each Netpage is printed with the Netpage logo at the bottom to indicate that it is a Netpage and therefore has interactive properties. The logo also acts as a “copy” button. In most cases pressing the logo produces a copy of the page. In the case of a form the button instead elicits a page giving the user the option to print the entire form document. And in the case of a secure document, such as a ticket or coupon, the button elicits an explanatory note or advertising page.

The default single-page copy function is handled directly by the relevant Netpage Page Server. Special copy functions are handled by linking the logo button to other URIs.

Once a Netpage form has been submitted, it is marked as submitted by the Netpage Page Server and cannot be submitted again. An attempt to do so elicits a status report indicating when it was submitted. A copy of the form can still be made, altered, and re-submitted.

5.9 THE HELP SYSTEM

The Netpage Printer has a single button labelled “help”. When pressed it elicits a single page of information. This information includes the following:

- status of printer connection
- status of printer consumables
- top-level help menu
- document function menu
- top-level Netpage Network directory

The help menu provides a hierarchical manual on how to use the Netpage System.

The document function menu includes the following functions:

- print a copy of a document

- print a clean copy of a form
- print the status of a document

A document function is initiated by simply pressing the button and then touching any page of the document. The status of a document indicates who published it and when, to whom it was delivered, and to whom and when it was subsequently submitted as a form.

The Netpage Network directory allows the user to navigate the hierarchy of publications and services on the network. As an alternative, the user can call the Netpage Network "900" number "yellow pages" and speak to a human operator. The operator can locate the desired document and route it to the user's printer. Depending on the document type, the publisher or the user pays the small "yellow pages" service fee.

The help page is obviously unavailable if the printer is unable to print. In this case the "error" light is lit and the user can request remote diagnosis over the network.

6 Personalized Publication Model

In the following discussion, news is used as a canonical publication example to illustrate personalization mechanisms in the Netpage System. Although news is often used in the limited sense of newspaper and newsmagazine news, the intended scope is wider.

In the Netpage System, the editorial content and the advertising content of a news publication are personalized using different mechanisms. The editorial content is personalized according to the reader's explicitly stated and implicitly captured interest profile. The advertising content is personalized according to the reader's locality and demographic.

6.1 EDITORIAL PERSONALIZATION

A subscriber can draw on two kinds of news sources: those that deliver news publications, and those that deliver news streams. While news publications are aggregated and edited by the publisher, news streams are aggregated either by a news publisher or by a specialized news aggregator. News publications typically correspond to traditional newspapers and newsmagazines, while news streams can be many and varied: a "raw" news feed from a news service, a cartoon strip, a freelance writer's column, a friend's bulletin board, or the reader's own e-mail.

The Netpage Publication Server supports the publication of edited news publications as well as the aggregation of multiple news streams. By handling the aggregation and hence the formatting of news streams selected directly by the reader, the server is able to place advertising on pages over which it otherwise has no editorial control.

The subscriber builds a daily newspaper by selecting one or more contributing news publications, and creating a personalized version of each. The resulting daily editions are printed and bound together into a single newspaper. The various members of a household typically express their different interests and tastes by selecting different daily publications and then customizing them.

For each publication, the reader optionally selects specific sections. Some sections appear daily, while others appear weekly. The daily sections available from The New York Times online, for example, include "Page One Plus", "National", "International", "Opinion", "Business", "Arts/Living", "Technology", and "Sports". The set of available sections is obviously specific to a publication, as is the default subset.

The reader extends the daily newspaper by creating custom sections, each one drawing on any number of news streams. Custom sections might be created for e-mail and friends' announcements ("Personal"), or for monitoring news feeds for specific topics ("Alerts" or "Clippings").

For each section, the reader optionally specifies its size, either qualitatively (e.g. short, medium, or long), or numerically (i.e. as a limit on its number of pages), and the desired proportion of advertising, either qualitatively (e.g. high, normal, low, none), or numerically (i.e. as a percentage).

The reader also optionally expresses a preference for a large number of shorter articles or a small number of longer articles. Each article is ideally written (or edited) in both short and long forms to support this preference.

An article may also be written (or edited) in different versions to match the expected sophistication of the reader, for example to provide children's and adults' versions. The appropriate version is selected according to the reader's age. The reader can specify a "reading age" which takes precedence over their biological age.

The articles which make up each section are selected and prioritized by the editors, and each is assigned a useful lifetime. By default they are delivered to all relevant subscribers, in priority order, subject to space constraints in the subscribers' editions.

In sections where it is appropriate, the reader may optionally enable collaborative filtering. This is then applied to articles which have a sufficiently long lifetime. Each article which qualifies for collaborative filtering is printed with rating buttons at the end of the article. The buttons can provide an easy choice (e.g. "liked" and "disliked"), making it more likely that readers will bother to rate the article.

Articles with high priorities and short lifetimes are therefore effectively considered essential reading by the editors and are delivered to most relevant subscribers.

The reader optionally specifies a serendipity factor, either qualitatively (e.g. do or don't surprise me), or numerically. A high serendipity factor lowers the threshold used for matching during collaborative filtering. A high factor makes it more likely that the corresponding section will be filled to the reader's specified capacity. A different serendipity factor can be specified for different days of the week.

The reader also optionally specifies topics of particular interest within a section, and this modifies the priorities assigned by the editors.

The speed of the reader's Internet connection affects the quality at which images can be delivered. The reader optionally specifies a preference for fewer images or smaller images or both. If the number or size of images is not reduced, then images may be delivered at lower quality (i.e. at lower resolution or with greater compression).

At a global level, the reader specifies how quantities, dates, times and monetary values are localized. This involves specifying whether units are imperial or metric, a local timezone and time format, and a local currency, and whether the localization consist of *in situ* translation or annotation. These preferences are derived from the reader's locality by default.

To reduce reading difficulties caused by poor eyesight, the reader optionally specifies a global preference for a larger presentation. Both text and images are scaled accordingly, and less information is accommodated on each page.

The language in which a news publication is published, and its corresponding text encoding, is a property of the publication and not a preference expressed by the user. However, the Netpage Network may provide automatic translation services in various guises.

6.2 ADVERTISING LOCALIZATION AND TARGETING

The personalization of the editorial content directly affects the advertising content, because advertising is typically placed to exploit the editorial context. Travel ads, for example, are more likely to appear in a travel section than elsewhere. The value of the editorial content to an advertiser (and therefore to the publisher) lies in its ability to attract large numbers of readers with the right demographics.

Effective advertising is placed on the basis of locality and demographics. Locality determines proximity to particular services, retailers etc., and particular interests and concerns associated with the local community and environment. Demographics determine general interests and preoccupations as well as likely spending patterns.

A news publisher's most profitable product is advertising "space", a multi-dimensional entity determined by the publication's geographic coverage, the size of its readership, its readership demographics, and the page area available for advertising.

In the Netpage System, the Netpage Publication Server computes the approximate multi-dimensional size of a publication's saleable advertising space on a per-section basis, taking into account the publication's geographic coverage, the section's readership, the size of each reader's section edition, each reader's advertising proportion, and each reader's demographic.

In comparison with other media, the Netpage System allows the advertising space to be defined in greater detail, and allows smaller pieces of it to be sold separately. It therefore allows it to be sold at closer to its true value.

For example, the same advertising "slot" can be sold in varying proportions to several advertisers, with individual readers' pages randomly receiving the advertisement of one advertiser or another, overall preserving the proportion of space sold to each advertiser.

The Netpage System allows advertising to be linked directly to detailed product information and online purchasing. It therefore raises the intrinsic value of the advertising space.

Because personalization and localization are handled automatically by Netpage Publication Servers, an advertising aggregator can provide arbitrarily broad coverage of both geography and demographics. The subsequent disaggregation is efficient because it is automatic. This makes it more cost-effective for publishers to deal with advertising aggregators than to directly capture advertising. Even though the advertising aggregator is taking a proportion of advertising revenue, publishers may find the change profit-neutral because of the greater efficiency of aggregation. The advertising aggregator acts as an intermediary between advertisers and publishers, and may place the same advertisement in multiple publications.

It is worth noting that ad placement in a Netpage publication can be more complex than ad placement in the publication's traditional counterpart, because the publication's advertising space is more complex. While ignoring the full complexities of negotiations between advertisers, advertising aggregators and publishers, it is clear that the Netpage System should ideally provide some automated support for these negotiations, including support for automated auctions of advertising space. Automation is particularly desirable for the placement of advertisements which generate small amounts of income, i.e. small or highly localized advertisements.

Once placement has been negotiated, the aggregator captures and edits the advertisement and records it on a Netpage Ad Server. Correspondingly, the publisher records the ad placement on the relevant Netpage Publication Server. When the Netpage Publication Server lays out each user's personalized publication, it picks the relevant advertisements from the Netpage Ad Server.

6.3 USER PROFILES

The personalization of news and other publications relies on an assortment of user-specific profile information:

- publication customizations
- collaborative filtering vectors
- contact details
- presentation preferences

The customization of a publication is typically publication-specific, and so the customization information is maintained by the relevant Netpage Publication Server.

A collaborative filtering vector consists of the user's ratings of a number of news items. As described in Section 4, it is used to correlate different users' interests for the purposes of making recommendations. Although there are benefits to maintaining a single collaborative filtering vector independently of any particular publication, there are two reasons why it is more practical to maintain a separate vector for each publication: there is likely to be more overlap between the vectors of subscribers to the same publication than between those of subscribers to different publications; and a publication is likely to want to present its users' collaborative filtering vectors as part of the value of its brand, not to be found elsewhere. Collaborative filtering vectors are therefore also maintained by the relevant Netpage Publication Server.

Contact details, including name, street address, ZIP Code, state, country, telephone numbers, etc., are by their nature global and are maintained by a Netpage Registration Server.

Presentation preferences, including those for quantities, dates and times discussed in Section 6.1, are likewise global and maintained in the same way.

The localization of advertising relies on the locality indicated in the user's contact details, while the targeting of advertising relies on personal information such as date of birth, gender, marital status, income, profession, education, etc., or qualitative derivatives such as age range and income range.

For those users who choose to reveal personal information for advertising purposes, the information is maintained by the relevant Netpage Registration Server. In the absence of such information, advertising can be targeted on the basis of the demographic associated with the user's ZIP or ZIP+4 Code.

Each user, pen, printer, publisher and publication is assigned its own globally unique identifier, and the Netpage Registration Server maintains the relationships between them.

Each user may be authorized to use any number of printers, and each printer may allow any number of users to use it. Each user has a single default printer, to which periodical publications are delivered. The server keeps track of which publishers a user has authorized print to the user's default printer.

Each user may have several pens, but a pen is specific to a single user. If a user is authorized to use a particular printer, then that printer recognizes any of the user's pens.

These relationships are illustrated in Figure 3.

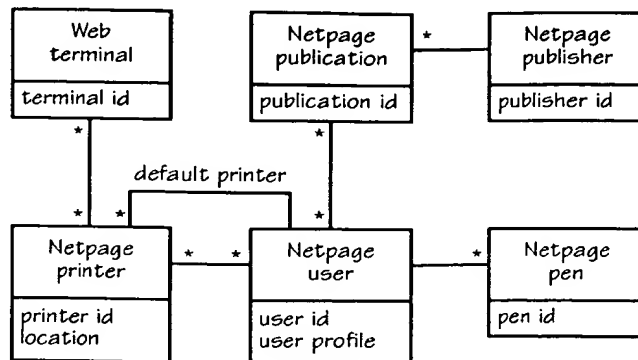


Figure 3. User registration relationships

The pen identifier is used, in the form of a URN, to locate the corresponding user profile maintained by a particular Netpage Registration Server, via the DNS in the usual way.

A Web terminal can be authorized to print on a particular Netpage Printer, allowing Web pages and Netpage documents encountered during browsing to be conveniently printed on the nearest Netpage Printer.

6.4 INTELLIGENT PAGE LAYOUT

The Netpage Publication Server automatically lays out the pages of each user's personalized publication on a section-by-section basis. Since most advertisements are in the form of pre-formatted rectangles, they are placed on the page before the editorial content.

The advertising ratio for a section can be achieved with wildly varying advertising ratios on individual pages within the section, and the ad layout algorithm exploits this. The algorithm attempts to co-locate closely tied editorial and advertising content, e.g. ads for roofing material placed specifically with the publication because of a special feature on do-it-yourself roofing repairs.

The editorial content selected for the user, i.e. text with associated images and graphics, is then laid out according to various aesthetic rules.

The entire process, including the selection of ads and the selection of editorial content, must be iterated once the layout has converged, to attempt to more closely achieve the user's stated section size preference. The section size preference can, however, be matched *on average* over time, allowing significant day-to-day variations.

6.5 DOCUMENT FORMAT

Once the document is laid out, it is encoded for efficient distribution and persistent storage on the Netpage Network.

The primary efficiency mechanism is the separation of information specific to a single user's edition and information shared between multiple users' editions. The specific information consists of the page layout. The shared information consists of the objects to which the page layout refers, including images, graphics, and pieces of text.

A text object contains fully-formatted text represented in the Extensible Markup Language (XML) [68] using the Extensible Stylesheet Language (XSL) [69]. XSL provides precise control over text formatting independently of the region into which the text is being set, which in this case is being provided by the layout. The text object contains embedded language codes to enable automatic translation, and embedded hyphenation hints to aid with paragraph formatting.

An image object encodes an image in the JPEG 2000 wavelet-based compressed image format [34]. The original DCT-based JPEG algorithm introduces negligible visual loss at compression ratios below 10:1 [65]. JPEG 2000 is planned to achieve the same quality at compression ratios 30% higher, i.e. at about 13:1 [35].

A graphic object encodes a 2D graphic in Scalable Vector Graphics (SVG) [71] format.

The layout itself consists of a series of placed image and graphic objects, linked textflow objects through which text objects flow, hyperlinks and input fields as described in Section 5.6, and watermark regions. These layout objects are summarized in Table 3. The layout uses a compact format suitable for efficient distribution and storage.

The layout is tagged with the version of the text-setting algorithm used by the Netpage Publication Server when the layout was first created, allowing the Netpage Printer to exactly reproduce the physical layout intended by the server.

Because Netpage Printer software is automatically upgraded over the Netpage Network, it is feasible for Netpage Printers to contain every version of the text-setting algorithm.

Table 3. Netpage layout objects

layout object	attribute	format of linked object
image	position	-
	image object URI	JPEG 2000
graphic	position	-
	graphic object URI	SVG
textflow	textflow id	-
	region ^a	-
	optional text object URI	XML/XSL
hyperlink	type	-
	region ^a	-
	handler URI	-
field	type	-
	meaning	-
	region ^a	-
watermark	region ^a	-

a. arbitrary multi-edged shape defined with spline paths

6.6 DOCUMENT DISTRIBUTION

As described above, for purposes of efficient distribution and persistent storage on the Netpage Network, a user-specific page layout is separated from the shared objects to which it refers.

When a subscribed publication is ready to be distributed, the Netpage Publication Server allocates, with the help of the Netpage Id Server, a globally unique id for each page, page instance, document, and document instance.

The server computes a set of optimized subsets of the shared content and creates a multicast channel for each subset, and then tags each user-specific layout with the names of the multicast channels which will carry the shared content used by that layout. The server then pointcasts each user's layouts to that user, and when the pointcasting is complete, multicasts the shared content on the specified channels. After receiving its pointcast, each Netpage Printer subscribes to the multicast channels specified in the page layouts. During the multicasts, each printer extracts from the multicast streams those objects referenced by its page layouts.

Once a printer has received all the objects to which its page layouts refer, the printer re-creates the fully-populated layout and then rasterizes and prints it.

The server also delivers each page layout to the relevant Netpage Page Server, which may be co-located with the Netpage Publication Server, or may be located elsewhere on the network. Thus the page layouts are persistently archived as Netpages. It is the responsibility of the Netpage Publication Server to preserve the shared objects referenced by the Netpages, to ensure that they are really persistent. It may choose to archive these shared objects elsewhere on the network at any time. The object URIs embedded in the Netpages allow the objects to move.

Under normal circumstances, the printer prints pages faster than they can be delivered. Assuming a quarter of each page is covered with images, the average page has a size of less than 400KB¹. The printer can therefore hold in excess of 100 such pages in its internal 64MB memory, allowing for temporary buffers etc. The printer prints at a rate of one page per second. This is equivalent to 400KB or about 3Mbit of page data per second, which is similar to the highest expected rate of page data delivery over a broadband network.

Even under abnormal circumstances, such as when the printer runs out of paper, it is likely that the user will be able to replenish the paper supply before the printer's 100-page internal storage capacity is exhausted.

However, if the printer's internal memory does fill up, then the printer will be unable to make use of a multicast when it first occurs. The Netpage Publication Server therefore allows printers to submit requests for re-multicasts. When a critical number of requests is received or a timeout occurs, the server re-multicasts the corresponding shared objects.

Once a document is printed, a Netpage Printer can produce an exact duplicate at any time by retrieving its page layouts from the relevant Netpage Page Server and retrieving the objects to which they refer from the network.

6.7 ON-DEMAND DOCUMENTS

When a Netpage document is requested ad hoc, it is personalized and delivered in much the same way as a periodical. However, since there is no shared content, delivery is made directly to the requesting printer, i.e. without the use of multicast.

1. 267 pixels per inch (ppi) 24-bit RGB, compressed using JPEG 2000 at a ratio of 13:1.

When a non-Netpage document is requested ad hoc, it is not personalized, and it is delivered via a designated Netpage Formatting Server which reformats it as a Netpage document. A Netpage Formatting Server is a special instance of a Netpage Publication Server. The Netpage Formatting Server has knowledge of myriad Internet document formats, including old favorites such as Adobe's Portable Document Format (PDF) [4], and Hypertext Markup Language (HTML) [70]. In the case of HTML, it makes use of the higher resolution of the printed page to present Web pages in a two-column format, with a table of contents and an index of links. By default it automatically includes all Web pages directly linked to the requested page. The user can tune this behavior via a preference.

The Netpage Formatting Server makes standard Netpage behavior, including interactivity and persistence, available on any Internet document, no matter what its origin and format. It hides knowledge of different document formats from both the Netpage Printer and the Netpage Page Server.

7 Security

7.1 CRYPTOGRAPHY

Cryptography is used to protect sensitive information, both in storage and in transit, and to authenticate parties to a transaction. There are two classes of cryptography in widespread use: secret-key cryptography and public-key cryptography. The Netpage Network uses both classes of cryptography.

Secret-key cryptography, also referred to as symmetric cryptography, uses the same key to encrypt and decrypt a message. Two parties wishing to exchange messages must first arrange to securely exchange the secret key.

Public-key cryptography, also referred to as asymmetric cryptography, uses two encryption keys. The two keys are mathematically related in such a way that any message encrypted using one key can only be decrypted using the other key. One of these keys is then published, while the other is kept private. The public key is used to encrypt any message intended for the holder of the private key. Once encrypted using the public key, a message can only be decrypted using the private key. Thus two parties can securely exchange messages without first having to exchange a secret key. To ensure that the private key is secure, it is normal for the holder of the private key to generate the key pair.

Public-key cryptography can be used to create a digital signature. If the holder of the private key creates a known hash of a message and then encrypts the hash using the private key, then anyone can verify that the encrypted hash constitutes the “signature” of the holder of the private key with respect to that particular message, simply by decrypting the encrypted hash using the public key and verifying the hash against the message. If the signature is appended to the message, then the recipient of the message can verify both that the message is genuine and that it has not been altered in transit.

To make public-key cryptography work, there has to be a way to distribute public keys which prevents impersonation. This is normally done using certificates and certificate authorities. A certificate authority is a trusted third party which authenticates the connection between a public key and someone’s identity. The certificate authority verifies the person’s identity by examining identity documents etc., and then creates and signs a digital certificate containing the person’s identity details and public key. Anyone who trusts the certificate authority can use the public key in the certificate with a high degree of certainty that it is genuine. They just have to verify that the certificate has indeed been signed by the certificate authority, whose public key is well-known.

In most transaction environments, public-key cryptography is only used to create digital signatures and to securely exchange secret session keys. Secret-key cryptography is used for all other purposes.

In the following discussion, when reference is made to the *secure* transmission of information between a Netpage Printer and a server, what actually happens is that the printer obtains the server’s certificate, authenticates it with reference to the certificate authority, uses the public key-exchange key in the certificate to exchange a secret session key with the server, and then uses the secret session key to encrypt the message data. A *session* key, by definition, can have an arbitrarily short lifetime.

7.2 NETPAGE PRINTER SECURITY

Each Netpage Printer is assigned a pair of unique identifiers at time of manufacture which are stored in read-only memory in the printer and in the Netpage Registration Server database. The first id is public and uniquely identifies the printer on the Netpage Network. The second id is secret and is used when the printer is first registered on the network.

When the printer connects to the Netpage Network for the first time after installation, it creates a signature public/private key pair. It transmits the secret id and the public key securely to the Netpage Registration Server. The server compares the secret id against the printer's secret id recorded in its database, and accepts the registration if the ids match. It then creates and signs a certificate containing the printer's public id and public signature key, and stores the certificate in the registration database.

The Netpage Registration Server acts as a certificate authority for Netpage Printers, since it has access to secret information allowing it to verify printer identity.

When a user subscribes to a publication, a record is created in the Netpage Registration Server database authorizing the publisher to print the publication to the user's default printer. Every document sent to a printer is addressed to a particular user and is signed by the publisher using the publisher's private signature key. The printer verifies, via the registration database, that the publisher is authorized to print the publication on the specified user's default printer and that the printer is the user's default printer. The printer verifies the signature using the publisher's public key, obtained from the publisher's certificate stored in the registration database.

The Netpage Registration Server accepts requests to add printing authorizations to the database, so long as those requests are initiated via a pen registered to the printer.

The user can authorize a Web terminal to print on a printer. This is useful if the user has a Web terminal in the home which is used to locate documents on the Web for printing. The one-time authorization proceeds as follows: the user prints a Web terminal authorization form. The Netpage Registration Server generates a short-lifetime one-time-use id for the Web terminal which is printed on the form, together with the URI of the printer. The Web terminal is used to navigate to a Netpage Registration Server registration site, where the one-time-use id is entered, as well as the URI of the printer. The Web terminal generates a signature public/private key pair. The server allocates a terminal id for the Web terminal, creates and signs a certificate containing the terminal id and the public key, and stores the certificate in the registration database. The URI of the printer, the Web terminal's terminal id, and the private signature key are stored locally on the Web terminal.

Whenever the Web terminal wishes to print on the printer, it sends the printer's designated Netpage Formatting Server a request containing the URI of the document to be printed, together with the terminal id, signed with the Web terminal's private signature key. On receipt of the request and before acting on it, the server verifies, via the registration server, that the terminal is authorized to print on the specified printer, and verifies the signature using the terminal's public key, obtained from the terminal's certificate stored in the registration database.

The user can print a list of current printing authorizations at any time, and revoke any which are being abused.

7.3 NETPAGE PEN SECURITY

Each Netpage Pen is assigned a unique identifier at time of manufacture which is stored in read-only memory in the pen and in the Netpage Registration Server database. The id uniquely identifies the pen on the Netpage Network.

A Netpage Pen can know a number of Netpage Printers, and a printer can know a number of pens. A pen communicates with a printer via a radio frequency signal whenever it is within range of the printer. Once a pen and printer are registered, they regularly exchange session keys. Whenever the pen transmits digital ink to the printer, the digital ink is always encrypted using the appropriate session key. Digital ink is never transmitted in the clear.

A pen stores a session key for every printer it knows, indexed by printer id, and a printer stores a session key for every pen it knows, indexed by pen id. Both have a large but finite storage capacity for session keys, and will forget a session key on a least-recently-used basis if necessary.

When a pen comes within range of a printer, the pen and printer discover whether they know each other. If they don't know each other, then the printer determines whether it is supposed to know the pen, i.e. because the pen belongs to a user who is registered to use the printer. If the printer is meant to know the pen but doesn't, then it initiates the automatic pen registration procedure. If the printer isn't meant to know the pen, then it agrees with the pen to ignore it until the pen is placed in a charging cup, at which time it initiates the registration procedure.

In addition to its public id, the pen contains a secret id and a secret key-exchange key, both intended for one-time use. These are also recorded in the Netpage Registration Server database at time of manufacture. During registration, the printer obtains the secret id from the pen. Because it is transmitted in the clear, it may be intercepted by someone listening in. The printer transmits the id securely to the Netpage Registration Server, which responds securely with the matching key-exchange key, together with a newly-generated secret id and key-exchange key. The printer generates a session key for the pen and transmits it to the pen encrypted using the one-time-use key-exchange key. It also securely transmits the new secret id and key-exchange key to the pen, which saves them for the next registration procedure. They now match the pen's record in the Netpage Registration Server database.

If the secret id transmitted in the clear from the pen to the printer is intercepted and used to retrieve the secret key-exchange key from the Netpage Registration Server before the printer queries the server, then the server rejects the printer's query because the secret id is out-of-date. Thus the printer knows that the pen has been compromised, and recommends that it be returned for repair.

When a previously unregistered pen is first registered, it is of limited use until it is linked to a user. A registered but "un-owned" pen is only allowed to be used to request and fill in Netpage user and pen registration forms [3], i.e. to register a new user to which the new pen is automatically linked, or to add a new pen to an existing user.

The pen uses secret-key rather than public-key encryption because of hardware performance constraints in the pen.

7.4 SECURE DOCUMENTS

The Netpage System supports the delivery of secure documents such as tickets and coupons. The Netpage Printer includes a facility to print watermarks, but will only do so on request from publishers who are suitably authorized. The publisher indicates its authority to print watermarks in its certificate, which the printer is able to authenticate.

The "watermark" printing process uses an alternative dither matrix in specified "watermark" regions of the page. Back-to-back pages contain mirror-image watermark regions which coincide when printed. The dither matrices used in odd and even pages' watermark regions are designed to produce an interference effect when the regions are viewed together - i.e. when looking *through* the printed sheet.

The effect is similar to a watermark in that it is not visible when looking at only one side of the page, and is lost when the page is copied by normal means.

As described in Section 5.8, pages of secure documents cannot be copied using the built-in Netpage copy mechanism. This extends to copying Netpages on Netpage-aware photocopiers.

Secure documents are typically generated as part of e-commerce transactions. They can therefore include the user's photograph which was captured when the user registered biometric information with the Netpage Registration Server, as described in Section 8.

When presented with a secure Netpage document, the recipient can verify its authenticity by requesting its status in the usual way. The unique id of a secure document is only valid for the lifetime of the document, and secure document ids are allocated non-contiguously to prevent their prediction by opportunistic forgers. A secure document verification pen can be developed with built-in feedback on verification failure, to support easy point-of-presentation document verification.

Clearly neither the watermark nor the user's photograph are secure in a cryptographic sense. They simply provide a significant obstacle to casual forgery. Online document verification, particularly using a verification pen, provides an added level of security where it is needed, but is still not entirely immune to forgeries.

7.5 NON-REPUDIATION

In the Netpage System, forms submitted by users are delivered reliably to forms handlers and are persistently archived on Netpage Page Servers. It is therefore impossible for recipients to repudiate delivery.

E-commerce payments made through the system, as described in Section 8, are also impossible for the payee to repudiate.

8 Electronic Commerce Model

8.1 SECURE ELECTRONIC TRANSACTION (SET)

The Netpage System uses the Secure Electronic Transaction (SET) [43] system as its payment system model. Although SET is not yet widely supported, it is comprehensive and elegant and will probably become dominant in the near future.

SET, having been developed by MasterCard and Visa, is organized around payment cards, and this is reflected in the terminology. However, much of the system is independent of the type of accounts being used.

In SET, cardholders and merchants register with a certificate authority and are issued with certificates containing their public signature keys. The certificate authority verifies a cardholder's registration details with the card issuer as appropriate, and verifies a merchant's registration details with the acquirer as appropriate. Cardholders and merchants store their respective private signature keys securely on their computers. During the payment process, these certificates are used to mutually authenticate a merchant and cardholder, and to authenticate them both to the payment gateway.

SET has not yet been adopted widely, partly because cardholder maintenance of keys and certificates is considered burdensome. Interim solutions which maintain cardholder keys and certificates on a server and give the cardholder access via a password have met with some success [14].

8.2 SET PAYMENTS

In the Netpage System the Netpage Registration Server acts as a proxy for the Netpage user (i.e. the cardholder) in SET payment transactions.

The Netpage System uses biometrics to authenticate the user and authorize SET payments. Because the system is pen-based, the biometric used is the user's on-line signature, consisting of time-varying pen position, tilt and pressure. A fingerprint biometric can also be used by designing a fingerprint sensor into the pen, although at a higher cost. The type of biometric used only affects the capture of the biometric, not the authorization aspects of the system.

The first step to being able to make SET payments is to register the user's biometric with the Netpage Registration Server. This is done in a controlled environment, for example a bank, where the biometric can be captured at the same time as the user's identity is verified. The biometric is captured and stored in the registration database, linked to the user's record and to the record of a particular Netpage Pen. The user's photograph is also optionally captured and linked to the record. The SET cardholder registration process is completed, and the resulting private signature key and certificate are stored in the database. The user's payment card information is also stored, giving the Netpage Registration Server enough information to act as the user's proxy in any SET payment transaction.

When the user eventually supplies the biometric to complete a payment, for example by signing a Netpage order form, the printer securely transmits the order information, the pen id and the biometric data to the Netpage Registration Server. The server verifies the biometric with respect to the user identified by the pen id, and from then on acts as the user's proxy in completing the SET payment transaction.

8.3 MICRO-PAYMENTS

The Netpage Network includes a mechanism for micro-payments, to allow the user to be conveniently charged for printing low-cost documents on demand and for copying copyright documents, and possibly also to allow the user to be reimbursed for expenses incurred in printing advertising material. The latter depends on the level of subsidy already provided to the user.

When the user registers for e-commerce, a network account is established which aggregates micro-payments. The user receives a statement on a regular basis, and can settle any outstanding debit balance using the standard payment mechanism.

The network account can be extended to aggregate subscription fees for periodicals, which would also otherwise be presented to the user in the form of individual statements.

8.4 TRANSACTIONS

Whenever a transaction originates through a Netpage form, the form handler has sufficient information, in the shape of the form's unique document instance id, to maintain transaction-specific state information. However, a transaction may also originate through a non-form page such as a printed catalog page, implying, for example, the existence of a virtual "shopping cart". In this case the relevant transaction state information is tied, indirectly, to the unique id of the user.

The Netpage Registration Server maintains an anonymous relationship between a user and a transaction handler via a uniquely numbered transaction, as illustrated in Figure 4. Whenever the user activates a hyperlink tagged with the "transaction" attribute, the Netpage Printer asks the Netpage Registration Server to translate the associated handler id, together with the pen id, into a transaction id. The transaction id is then submitted to the hyperlink transaction handler. For efficiency, the printer caches transaction ids.

The transaction handler maintains state information indexed by transaction id. It is able to retrieve user-specific state information without explicit knowledge of the user.




Figure 4. User transaction

APPLICATIONS AND BUSINESS MODELS

9 Applications

The Netpage Network has the potential to subsume a wide variety of applications in both traditional and electronic media. This section sketches the following possibilities:

- personalized subscriptions to newspapers, magazines and comics
- subscriptions to freelance columns and bulletin boards
- on-demand newspapers, magazines and comics
- on-demand flyers and product brochures
- on-demand books
- television infotainment printing
- e-commerce purchasing from online and traditional catalogs
- delivery of statement and invoices, with online payment
- delivery of secure document (tickets, coupons and licenses)
- perfect copying with copyright micro-payments
- mail replacement
- delivery of greeting cards
- form printing, fill-in, and submission
- delivery of e-mail and facsimile
- active business cards
- on-demand document delivery on corporate intranets
- provision of government services
- Web browsing, searching and printing
- photo album creation
- persistent searchable note taking
- exam taking and homework
- computer printing

Many of these applications are described in detail in [3].

9.1 PERSONALIZED SUBSCRIPTIONS

The strength of the Netpage Network lies in automatically delivering subscribed periodicals, at a print quality equalling or exceeding that of their traditional counterparts, with editorial content personalized to individual interests, advertising content localized and targeted to individual localities and demographics, and advertising directly linked to detailed product information and product purchasing.

9.1.1 Newspapers and Magazines

The Netpage Network offers a new delivery mechanism to the \$24 billion newspaper and newsmagazine market which is both more cost-effective than centralized printing and distribution, and allows more fine-grained targeting of advertising.

The simplest form of news personalization involves selecting a news publication and choosing which daily and weekly sections to receive. The simplest form of advertising personalization is tuned to the demographics associated with the subscriber's ZIP or ZIP+4 code. Even with these entry-level forms of personalization, the Netpage Network offers a compelling distribution model for news.

Users receive paper publications with the usability of their traditional counterparts, but with interactive properties. At the press of an on-page button, a user can print an article giving the background to a news story, print a personalized product brochure, or add a product to the virtual shopping basket.

Any magazine normally printed on lightweight paper stock is equally well-suited to distribution on the Netpage Network. However, since Netpage Printers don't carry heavier paper stocks and don't provide wrap-around binding, the Netpage Network is less well-suited to the distribution of so-called "glossy" magazines.

9.1.2 Freelance Columns and Bulletin Boards

Users can choose to subscribe to individual freelance columns, cartoons, etc. These can be integrated into a user's daily news document, or printed individually. Freelancers can choose to receive micro-payments from their subscribers, freeing them from maintaining their own subscriber databases. The Netpage Network provides mechanisms for handling micro-payments.

Users can also subscribe to the "bulletin boards" of friends: collections of news, announcements, pictures etc., which work much like freelance columns.

9.2 ON-DEMAND PUBLICATIONS

The Netpage Network can deliver, on demand, current and back issues of periodicals normally delivered on subscription, including newspapers, magazines, and comics. To maintain the interactivity of all Netpages ever printed, the Netpage Network keeps all published content online at all times. Unlike the Web, where hyperlinks become unreliable over time, content on the Netpage Network never expires.

9.2.1 Flyers and Product Brochures

The Netpage Network makes high-quality flyers and product brochures instantly available, linked to advertisements and entries in printed catalogs.

Brochures are always up-to-date, and link to e-commerce, e-mail, and automatic telephone call-back. Brochure links can provide "click-through" fees to linking documents, and subsidized printing to users.

Active Netpage advertisements can appear in printed paper publications sold or delivered through traditional outlets, including newspapers and magazines sold on news stands, posters appearing in public places (including in buses, trains and taxis), and books sold in bookshops. Buying guides and travel guide books are particularly suited to containing active Netpage advertising.

9.2.2 Books

Users can obtain the latest best-sellers or rare "out-of-print" (a soon-to-obsoleted term) editions on demand, printed in column format with a text size chosen by the user. A typical 300-page paperback fits on as little as 40 sheets of Letter paper. Slip-on covers are available for robust handling.

Titles which have outlived their copyright period are available for free. Other titles are heavily discounted for Netpage delivery, since publishers avoid the costs of printing, inventory storage, and delivery.

Colorful children's books reproduce immaculately. When they've been loved to death, they can be printed again, and again.

Children's coloring-in books and puzzles are available just when they're needed on a rainy day.

9.2.3 Television Infotainment

Users watching infotainment programs on television can print the associated informational material on their Netpage Printer by pressing the "print" button on their remote control at the appropriate time. The material may be a report on a new medical procedure, the plans for a do-it-yourself bookshelf, or a list of top investment opportunities. It may also be a subscription form for the print publication the program is promoting.

The television and/or remote control is suitably modified or augmented to route the print request to the Netpage Printer. The television or remote control may, for example, include a radio transmitter or transceiver which allows it to communicate with the Netpage Printer. When the print button is pressed, a message is sent to the Netpage Printer which contains the identity of the television channel. The television channel id, date and time together uniquely identify the Netpage document associated with the program at that time. The printer constructs a URN based on the television channel id and determines the address of the television channel's document server via the DNS in the usual way. It then retrieves the document from the document server based on date and time.

A potentially more flexible approach embeds the URN of the document in the television signal itself, for example as a closed caption. In analog television signals there is plenty of room for information in the vertical blanking interval (VBI). The advantage of leveraging support for closed captioning, in particular, is twofold. Closed caption decoders are now standard components in normal-sized television sets, and closed captioning is fully supported in digital television standards that are now being deployed. Other VBI-based schemes require custom decoders and may not survive the transition to digital television.

9.3 E-COMMERCE

9.3.1 Online Purchases

The Netpage Network supports a similar level of online purchasing as the Web, but in paper-based medium which presents like a high-quality printed catalog.

A user can navigate the retailer's online Netpage catalog, printing catalog pages as they're needed and adding items to a virtual shopping trolley. The contents of the shopping cart can be listed at any time, and items can be struck from the list at the stroke of the pen. Pressing a "proceed to checkout" button at any time elicits a completed order form just waiting for the user's signature. The payment card account number is securely shown in the usual 1234 56** **** *789 format. The user's signature authorizes the payment.

9.3.2 Catalog Purchases

Rather than buying from an online Netpage catalog, the user can select items from a traditionally-printed catalog which contains active Netpage links.

9.3.3 Statements and Invoices

Statements and invoices can be securely and auditably delivered, and can be automatically filled in with the user's default payment details without the sender knowing those details.

The user's signature can authorize the payment as normal.

9.3.4 Secure Documents

Retailers can securely issue tickets and coupons over the Netpage Network, printed with difficult-to-forgo watermarks.

Agencies of various kinds can issue licenses printed with watermarks and the user's own photograph.

As described earlier, recipients can verify the authenticity of secure documents using a standard Netpage Printer or a special verification pen.

9.3.5 Copyright Copying

Any printed version of a Netpage document becomes an easy means to printing another perfect copy. When a copy is made, the Netpage Network can automatically transfer a micro-payment from the copier to the copyright holder.

Trivial copyright fees are universally respected but seldom paid because of the inconvenience. The Netpage Network offers micro-payment convenience and the quality of an original copy.

9.3.6 Authenticated Signing and Exchange of Legal Documents

Netpage supports authenticated delivery and printing of documents, and automatically captures and verifies signatures entered in signature fields. It can therefore support authenticated delivery and signing of legal documents.

9.3.7 Retail Signature Capture

Point-of-purchase signature capture is a common requirement when the customer uses a payment card or charge card. Some retailers are now adopting digitizing tablets with LCDs to capture signatures, both to reduce paperwork and potentially to implement signature verification in the future. A Netpage paper roll which is printed with the details of each transaction and then captures the customer's signature provides an alternative. The customer benefits from the ease of using an ink pen on paper, while the retailer benefits from online signature capture and verification.

9.4 COMMUNICATION

9.4.1 Mail

The Netpage Network, once widely subscribed, can be used to deliver numerous instances of regular mail-outs, particularly statements and invoices as discussed in Section 9.3.3. As already described, the Netpage System only delivers mail from sources authorized by the recipient.

The United States Postal Service delivers 107,000,000,000 pieces of First Class mail each year [65], a large number of which are regular in nature.

9.4.2 Greeting Cards

A user can select a greeting card from an online catalog, add a handwritten message, and dispatch it via the Netpage Network. Cards can be addressed to other Netpage users, and to normal postal addresses. In the latter case the card is printed at the service center closest to the recipient, automatically placed in an envelope, and mailed through the local mail system.

Netpage users can choose to receive cards from anyone, or only from authorized friends.

9.4.3 Forms

Forms of all kinds can be printed on the Netpage Printer, filled in by hand, and submitted directly over the Netpage Network. Submission is secure and cannot be repudiated.

Handwriting is automatically recognized by the system. The digital ink of the handwriting is attached to the form in case a human clerk needs to re-interpret the handwriting. Automatic "handwriting bots" on the network can assist with the recognition task, automatically giving the user semi-intelligent feedback to elicit disambiguation.

Any interactive Netpage "application", including e-commerce and e-mail, uses forms of various kinds.

9.4.4 E-Mail

E-mail forms can be printed on demand and filled in by hand. The handwritten name of the recipient is converted to facilitate delivery, but the rest of the message is delivered as digital ink, just as the user intended. If the recipient is computer-based rather than Netpage-based, all of the handwriting can be automatically converted, with the digital ink sent as an attachment (since it may contain hand-drawn diagrams etc).

Netpage users can choose to receive e-mail from anyone, or only from authorized friends.

Each user maintains a list of contacts which allows e-mail to be addressed by name. Users never deal with Netpage user ids directly. If the specified name is ambiguous, the system prints a list of choices. The use of nicknames helps to avoid ambiguity.

9.4.5 Business Cards

Each user's business card can act as a convenient one-time e-mail authorization token. It is encoded as a Netpage and contains a button which, when pressed by the card recipient, adds the card owner to the recipient's contact list and authorizes the recipient to send the card owner e-mail.

The business card can also contain buttons which link to company information, personal information, and even product information if the user is a salesperson.

The Netpage Network provides a bureau service which allows users to order active Netpage business cards.

9.4.6 Facsimile

Facsimile forms can be printed on demand and filled in by hand. The handwritten telephone number is converted to facilitate delivery, but the rest of the message is delivered as bitmapped digital ink, just as the user intended.

9.4.7 Conferencing

Geographically distributed participants can communicate graphical ideas via a shared Netpage "whiteboard" while engaged in a telephone conference. Every time a participant adds a modification or a new diagram to the whiteboard the other participants are given a new copy of the page. New pages can be started at any time, but older pages or older versions of pages can still be marked up and re-distributed. Each participant ends up with a full history of the interaction. Multiple colors and line styles can be chosen from a palette.

Each printed page is timestamped and marked with the name of the author of the latest change. The latest change can be shown in a different style to the older content, to indicate what is new. For example, if the whiteboard is monochrome, then the latest change can be shown in color, or if the whiteboard is in color, then the latest change can be shown in bold strokes.

9.4.8 Collaborative Document Markup

Geographically distributed participants can mark up a shared Netpage document while engaged in a telephone conference. Every time a participant adds a modification or a new diagram to the document the other participants can be given a new copy of the page. A participant chooses to share their contribution with the other participants by pressing the <update> button which appears on every page. Older versions of pages can be marked up and re-distributed at any time. Each participant ends up with a full history of the interaction. Each participant's contribution can be shown in a different color, so that contributions can be distinguished. A key printed on each page or an initial status page indicates which color corresponds to which participant.

Each printed page is timestamped and marked with the name of the author of the latest change. Beyond color-coding each user's contribution, the latest change can be shown in a different style to older markup, to indicate what is new. For example, the latest change can be shown in bold strokes.

A session chairperson creates the markup session and links a Netpage document to it. The chairperson then e-mails session invitations to potential participants. Each invitation contains an <accept> button, which, when pressed, adds the invitee to the session. The invitation also typically contains the date and time that the session is intended to start, presented in the local time of the invitee, as well as the name of the document and the purpose of the markup. The chairperson receives e-mail every time an invitee accepts the invitation. The chairperson can distribute a markup session control page to the participants at any time, either before or during the actual session. The control page contains the session details, lists the session participants and shows their assigned markup colors, and contains a link to the document being marked up whereby the participant can obtain a copy. It also contains an <update> button which is used by the participant during the session to distribute their contribution to the other participants.

The chairperson's session control page also contains <start session> and <end session> buttons. No markup is possible before the session is started or after the session is ended.

The chairperson can designate non-contributory participants - i.e. ones who can monitor the markup session but cannot mark up the document themselves.

The session invitation also contains details on how to establish the telephone conference which accompanies the session. Where possible, the Netpage System controls the telephone system to initiate the conference.

A non-Netpage document can be marked up by first turning it into a Netpage document - i.e. by printing it to a Netpage Printer in Netpage mode (as described in Section 9.6.5).

9.4.9 Postage Stamps

Postage stamps can now be purchased over the Internet and printed on an inkjet printer. Netpage provides an even more compelling purchase and delivery mechanism for postage stamps because of its seamless payment mechanisms (for explicit payments and implicit micro-payments) and its support for secure documents.

As Netpages, Netpage-delivered postage stamps can contain both attractive graphics and invisible machine-readable information such as source and destination codes, allowing mail systems to track mail items and automatically route them.

9.4.10 Voicemail

The Netpage System can provide access to nominated voicemail mailboxes maintained by a telephone system. The help page can contain a <voicemail> button which prints a list of voicemail messages received by the user of the pen. Each message shows the date and time of receipt, the name of the caller (determined by looking up an address book using the phone number provided by caller ID), and the status of the message - i.e. whether it has been played back yet. The user can organize the message list into a hierarchy of folders as described in Section 6 in [3]. A message can be printed, in which case the system attempts to convert the message speech into text. The message can be played back through an optional speaker in the Netpage Printer. The message can also be played back through the user's mobile phone, either by the Netpage Printer controlling the mobile phone via its short-range radio interface (i.e. Bluetooth), or by the Netpage Printer instructing the telephone system to call the mobile phone and automatically play back the specified message.

9.4.11 Interactive Plans

Interactive plans can be printed as Netpages. The plans can be linked to a database and can be marked up with a Netpage Pen. Multiple people can collaborate on a single paper plan, and small changes can be made to a plan on the shop floor, without requiring access to a computer or an operator with computer skills.

If the plan is of a three-dimensional object, then different views of the object can be selected using a 3D view control, as described in Section 9 in [3].

The plan might be for a shoe design, and might show a dimensioned view of the shoe as well as a perspective view from a selected position. New dimensions can be entered in numeric fields and saved to the same plan file or to a new plan file.

9.5 CORPORATE AND GOVERNMENT

9.5.1 Corporate Intranets

An organization can use a private intranet-based Netpage Network to implement a document repository and efficiently distribute documents on demand.

9.5.2 Government Services

Government can provide access to services via the Netpage Network. The network can obviate the need to visit government offices to obtain forms and submit forms, and the network can be used to efficiently deliver the results of submissions.

9.6 PERSONAL

9.6.1 Web Browsing, Searching and Printing

Users can browse the World Wide Web via their Netpage Printer using paper and pen as the user interface. Netpage forms can provide emulation of HTML forms. Only dynamic media objects may fail to print meaningfully.

A Netpage Printer can be the ideal output device for documents encountered while browsing the Web, whether the browsing is terminal-based or Netpage-based. An increasing number of print-ready documents are being published on the Web.

Entire Web sites can be compactly formatted for print, since a printed page has a much greater information-bearing capacity than a computer screen, and perused in a more leisurely fashion. The Netpage Formatting Server automatically adds a table of contents and an index of referenced pages.

Links on Web pages remain active when printed on Netpages. Clicking on a link causes the linked Web page to be printed.

Web pages and search results are printed on both sides of the paper, allowing one sheet of paper to hold the equivalent of ten SVGA screens of information. Thus the user sees 200 links on paper, compared with only 20 on a computer screen.

9.6.2 Photo Album

A Netpage Printer can be enhanced with an infrared data connection (IrDA) to allow it to accept images wirelessly from a digital camera. Alternatively, digital cameras with USB ports can be directly connected. Images can be automatically archived on the Netpage Network, and individual photos, both regular and poster-sized, can be printed at photo quality on the printer.

Smart layout software can assist with the interactive creation of photo album pages, ready for insertion in a cumulative family album.

9.6.3 Notebook

A pre-printed Netpage notebook can be used to capture handwritten notes in a persistent fashion. Notes are captured on the network and are optionally timestamped and signed by a certificate authority to provide undeniable proof of priority, for example with respect to a patent application. Since the handwriting is recognized and converted to text for indexing purposes, the notes can also be searched by keyword, date, etc.

9.6.4 Examination Papers and Homework

Students can sit exams in a controlled manner while physically distant from the classroom.

An examination paper can be delivered at a specific time to a distributed collection of students. The students can then be required to answer the questions on the paper within the allotted time. Their answers can be automatically captured at the end of the allotted time using the "blackboard" form model, and in many cases automatically graded. Once again, unrecognized input can be routed to a human operator.

Multiple choice questions can be automatically graded.

Homework in general can be distributed and captured in a similar way.

9.6.5 Computer Printing

A Netpage Printer can be the ideal output device for a personal computer or workstation, whether directly-connected or on a local-area network. The speed of the Netpage Printer and the quality of its output make it attractive to existing computer users.

The Netpage Printers have a USB port as a standard feature.

A Netpage Printer can be used to print output from a computer system both as normal pages and as Netpages. When printed as Netpages, the pages are implicitly stored persistently on the Netpage Network, and the printed Netpages therefore have much of the interactivity of Netpages which originate on the network.

9.6.6 Netpage Photocopying

A Netpage-aware photocopier can be used to turn normal pages into Netpages in a similar way. When a normal (non-Netpage) page is copied, a digital version of the page is stored persistently on a Netpage Page Server, and the page copied is printed as a Netpage which is linked in the usual way to its digital counterpart. The digital version of the page may consist purely of an image of the page, or it may include text obtained via optical character recognition, thus allowing text-based searching.

When a Netpage is copied, a perfect digital copy is produced, subject to the usual Netpage security considerations as discussed in Section 7.4.

9.6.7 Games of Chance

Many gambling games can be delivered through Netpage. With instant lotto, just print the form, mark the appropriate numbers, and press the submit button. Win or lose, the Netpage response is instant.

With a group lottery or Keno, again just print the form, mark the appropriate numbers, and press the submit button. Wait for the announcement of the winning numbers, for example on television. Winners are also notified by Netpage.

9.6.8 Multi-Player Gaming

Multi-player games can be standard games such as chess. The players take turns marking moves on a Netpage showing the status of the game. After each move the player presses the submit button. The update is automatically printed on the other player's Netpage Printer. The opposing player can of course also be a game-playing program running on the Netpage Network.

An exploratory multi-player game (such as Ultima Online) can be presented as a Netpage. The Netpage shows the current view, the collection of items currently held, a map of the

area explored so far, and the chatter of nearby players. The form provides a text input field for chatting with other players, and command fields for navigating the game world. The submit button submits any chat text and navigation commands and produces an updated page. Previously printed pages form part of an evolving map of the game world, and provide the basis for hyperjumping to different locations in the game world under the appropriate circumstances.

9.7 CONSUMER ELECTRONICS

9.7.1 Netpage Digital Camera

A Netpage Printer built into a digital camera is used both as a Netpage Printer and as a photo printer. A digital camera containing a Netpage Printer can provide the option of printing its photos as Netpages, thus providing persistent storage of digital negatives on the Netpage Network. Each printed photo allows a pristine copy of itself to be produced at any time, even if the original print is severely degraded.

The digital negative also records the date, time and location the photo was taken, and the settings of the camera at the time. It also records any artistic effect applied to the photo when it was printed [1], so that the actual output can be reproduced.

The camera user can route photos to particular personal albums on the Netpage network. The user can grant other Netpage users (such as family members) access to particular albums. People can thus easily share photos without having to get together and swap negatives or prints. This may be particularly convenient when the camera user is travelling. The camera user may also choose to defer printing of photos while travelling, and instead retrieve and print the photos on returning home, possibly in the form of printed pages suitable for inclusion in a normal photo album (see Section 9.6.2).

9.8 ADVERTISING

9.8.1 Reverse-side Advertising

In many highly interactive Netpage applications, much of the information is supplied to the user one printed sheet at a time, since this is the minimum unit of Netpage output. Although a printed sheet has the capacity to hold two printed Netpages, i.e. one on each side of the sheet, there is in many cases insufficient information to fill two complete Netpages. When a sheet is printed with free space on the reverse side, unobtrusive advertising can be inserted. Although the immediate impact of the advertising is reduced because the advertising doesn't interrupt the user, the advertisement has persistence, and can be encountered some time after the interactive session if the user turns over the sheet while tidying up or archiving.

The Netpage system detects sheets which are to be printed on one side only, and automatically prints selected advertising material on the reverse side.

Although each user may have the ability to disable reverse-side printing of advertising, the Netpage system is likely to reward users for accepting advertising, since advertising and concomitant e-commerce provides the basis for subsidizing users' running costs.

9.9 PRODUCT INFORMATION AND REGISTRATION

9.9.1 Product Information

The packaging of a product can be printed as one or more Netpages containing buttons which link to various kinds of information about the product. A milk carton might contain buttons linked to information about the manufacturer, the product, the product in detail, recipes which involve the product, and advertising promotions involving the manufacturer or product, such as competitions. A music compact disc might contain buttons which link to information about the band and about the songs on the disc itself. A consumer product such as a television might contain a button which links to usage instructions.

In each case the consumer presses the button with a Netpage pen and receives the linked information on the Netpage Printer. Information is always delivered in the language of the user. This is particularly useful in the case of consumer electronics products which are sold in international markets.

Products can indicate both their instance identity (serial number) and class identity (universal product code). Consumers can thereby obtain instance-specific information from the Netpage Network, such as known problems, recalls, use-by-date, etc., as well as class-specific information, such as generic product information.

9.9.2 Product Registration

Further to Section 9.9.1, a product can also contain a <register> button which, when pressed with a Netpage Pen, registers the Netpage user as the owner of the product. The user's contact information, which is already recorded on the Netpage System, is automatically transmitted to the product manufacturer who can record it in their customer database. The registration process automatically adds the manufacturer to the user's e-mail contact list, thus allowing the manufacturer to send the user e-mail relevant to the product, such as related special offers, recall notices, etc. If the manufacturer abuses their e-mail privileges, the user can bar them in the usual way.

9.9.3 Product Information via Universal Product Code Bar Code

Many of the advantages of printing or affixing Netpages on the surfaces of products (as described in Section 9.9.1 above) can be gained by simply enhancing the Netpage Pen to decode the universal product code (UPC) bar codes which universally adorn products. The pen transmits the UPC to the Netpage Printer, tagged as a UPC rather than as digital ink. The printer constructs a UPC-specific URI and locates the corresponding resource on the Netpage Network via the DNS in the usual way.

The request thus ultimately elicits printed product information. The product page can also include the facility to enter the serial number of the device and register the user's ownership of it via a <register> button, with all the advantages described in Section 9.9.2 above.

Product manufacturers can thus gain the benefits of Netpage linking for their entire installed base of products without making alterations to the products themselves.

9.9.4 Product Information via Help Button

A physical help button can be installed in consumer devices such as televisions, kitchen appliances, etc.. The help button is coupled to a radio transceiver in the device which communicates a information request to the nearest Netpage Printer. The request can consist of

just the universal product code (UPC) of the device, or can contain instance-specific information such as the serial number of the device as well.

The information request is handled by the Netpage Printer in the same way as described in Section 9.9.1 or Section 9.9.3 above, and ultimately elicits printed product information.

9.9.5 Context-Specific Product Help

Further to Section 9.9.1, the entire surface of a product can be printed as a Netpage. Pressing on any part of the surface with a Netpage Pen can then elicit printed Netpage output which provides product-specific help. The help is either specific to the area pressed, or relates to the product as a whole. Thus the user of the product has instant access to helpful information about specific features of a product as well as the product as a whole. Each feature-specific help page is linked to the entire product manual.

9.9.6 Product Ownership Tracking

Further to Section 9.9.2 and Section 9.9.5, the entire surface of a product can be printed as a Netpage. Pressing on any part of the surface with a Netpage Pen can then elicit printed Netpage output which describes the product and its current ownership.

After the product is purchased, pressing on any part of the surface automatically registers the product in the name of the owner of the Netpage Pen.

Anyone can determine the ownership of a product offered for sale simply by pressing on any part of its surface with a Netpage Pen.

Ownership can only be registered by a new owner if the current owner has relinquished ownership by signing the the "sell" portion of the product's status page. This places the product in an "un-owned" state.

Product information and ownership is maintained either by the product manufacturer, as a service to its customers, or by a profit-oriented third party. A manufacturer-maintained product registration database has many advantages (see Section 9.9.2 above), but has the disadvantage that it may be abused for advertising purposes. A third-party-maintained product registration database has the disadvantage that it may incur a usage fee.

The shipping computer system of a product manufacturer can automatically transfer ownership of products from the manufacturer to the distributor or retailer, and so on down through the pre-sale chain. The retail computer system of the retailer can automatically mark each sold item as free, or transfer ownership directly to the holder of the payment card used to pay for the product. The customer can also use a Netpage Pen at the point of sale to register immediate ownership of the product.

Traditional clearing-houses for stolen goods, such as pawn shops, can be required by law to check the ownership of all products presented to them. Since a Netpage-encoded product has an invisible encoding on most or all of its surface, it is difficult for a thief to remove it or even tell if it has been successfully removed. Conversely, it is incumbent on a potential buyer of a product to ensure that a clean reading can be obtained from its surface so that its ownership can be indisputably established.

Where a product is leased or otherwise subject to complex or multiple ownership, the product registration database can reflect this and thus alert a potential buyer.

9.10 GEOGRAPHIC INFORMATION

9.10.1 Route Planning

An automobile navigation system typically determines the location of the car via GPS, shows a local map to the driver on a color display, assist the driver in planning a route, and gives the driver verbal instructions on following a route, cued to the changing location of the car. The Netpage System provides a convenient way for a driver to specify a destination to the navigation system's route planner. The user can press on a location on a Netpage street map and then press the <destination> button on the navigation system. The navigation system automatically programs the map location as the destination. The user also can press the <location> button on an advertisement or on someone's business card and then press the <destination> button on the navigation system. This yields a potentially more accurate location than one derived from a map.

The same technique allows a location to be specified as the destination to any route planning system, for example one provided by the Netpage System itself. A Netpage route planning system accepts a destination location in the same way. It can assume that the starting location is the location of the printer through which the user is interacting, but allows the user to explicitly specify a different starting location, as well as to modify the default route and specify route constraints in the usual way supported by route-planning systems.

9.10.2 Globe

The surface of a globe, i.e. a sphere representing the earth, can be printed as a Netpage. A user can then interact with the globe via the Netpage Pen. The user can request information about geographical features, cities, countries etc.; can request more detailed maps; can request the distance between two locations; can request a route including a number of locations; etc. In each case the Netpage Printer responds with the printed information.

In the case of detailed maps, successively more detailed maps can be requested, culminating in street maps where available. Many different map views are possible, e.g. showing topography, vegetation, annual rainfall, seasonal temperature, population, malaria incidence, etc. Multiple views can also be combined.

The user can select a number of points of the globe simply by clicking on them in succession with the pen. The user can also select a geographic region by circumscribing it with the pen. Subsequent operations, selected by pressing Netpage buttons printed on the base of the globe or on a handy command card, act on the selected points or the selected region. Different command cards can be provided for users of different sophistication, e.g. children and adults.

Ideally the globe rolls freely on its base, rather than rotating on an axis, so that *any* area on its surface can be rotated into view. A cupped base may be inset with ball bearings, for example. The surface of the globe is coated so that extended interaction with the base, as well as with the Netpage Pen, does not degrade the Netpage surface encoding.

The globe can provide audio output as an alternative to printed output. This may be particularly appropriate for interactive guessing games. The globe can ask the user to select the country which has a certain set of attributes or contains a certain geographic feature, and can rate the choice or give a hint once the user makes a choice. Guessing games can in general involve any number of players, or just one player and the globe.

10 Business Models

10.1 SYSTEM PRINCIPLES

The Netpage Network leverages the open technology and extensive infrastructure of the Internet. The widespread acceptance and growth of the Netpage Network is predicated on open competition rather than monopolistic practices.

However, to provide an incentive to early investors, semi-exclusive licenses to Mem-jet-based Netpage Printer designs will be offered, as well as licenses to manufacture paper and ink consumables.




Figure 5. Open structure of Internet-based Netpage Network

The network supports any number of independent participants, some of which have complementary roles, and some of which compete. The open structure of the network is illustrated in Figure 5. Content-related participants include the following:

- news sources
- publishers
- news aggregators
- freelance artists, writers, cartoonists
- direct mailers
- advertisers
- advertising aggregators
- banks
- merchants

Infrastructure-related participants include the following:

- server suppliers
- network storage providers
- communications carriers
- Internet service providers (ISPs)
- printer dealers
- printer installation and servicing companies
- ink and paper consumables dealers
- consumables delivery companies

Technology-related participants include the following:

- research and development
- chip makers (printheads, controllers, QA)
- printer manufacturers
- ink and paper consumables manufacturers

The strength of the network lies in the fact that publication and delivery are completely decoupled. This allows the delivery infrastructure to grow independently of the participation of publishers.

10.2 BOOTSTRAPPING

Because consumers are unlikely to be motivated to acquire a Netpage Printer until a variety of publications and services are available, and because publishers will wait for an installed base before participating, the key to bootstrapping the network is to bundle the printer with a publication or service subscription, and possibly trimming profit margins in the growth stage.

There are several ways the manufacturing cost of the Netpage Printer, assumed to be well below \$100, can be subsidized. Printer-based distribution can eliminate existing distribution costs, offsetting the printer cost. The printer can provide a new mechanism for delivering advertising, with advertising profits offsetting the printer cost. And the printer cost can be built into the subscription fee for a publication or service.

The cost of printing and delivering a newspaper normally exceeds the price of subscription [27,51]. The real profit lies in the advertising. The cost of the Netpage Printer is easily exceeded by one year's cost savings, allowing a Netpage subscription, including a "free" printer, to be priced lower than a traditional subscription. A Netpage subscription offered to a customer already on the network would be priced correspondingly lower still.

If the publication or service delivered via the Netpage Printer is sufficiently lucrative, then the publisher or provider may be able to subsidize not only the printer itself, but also its running costs. This can include Internet access, paper and ink consumables, and servicing. Demographics-informed advertising may fall into this category. The more information customers reveal about themselves, the greater the value of the advertising to the advertisers, and so the greater the level of subsidization that can take place.

Early investors who subsidize the installation of Netpage Printers may be able to recover the investment and turn a profit merely by charging other publishers an access fee to the

printers they “own”, perhaps for an interim period after installation, according to a network-wide agreement. They may also be able to earn commissions on click-throughs and e-commerce transactions originating on pages printed on “their” printers.

Similar approaches are already emerging in the general Internet market. In the “FreePC” and related models [37], personal computers are bundled with Internet access, and the whole package is fully or partially subsidized by advertising and e-commerce.

Most content-related participants in the Netpage Network, and even Internet service providers, can benefit from directly investing in Netpage Printer deployment.

10.3 MATURITY

Many of the bundling approaches are likely to remain applicable once the network becomes widely subscribed. It is possible that the bundling of the appliance (i.e. the Netpage Printer) with the service (be it Internet access or a publication subscription) will remain the dominant means of distributing the appliance, as it is in the cellular telephone market, and as an increasing number of companies, including IBM [18], are beginning to believe it should be in the personal computer market.

10.3.1 News Publishers and News Aggregators

News publishers with strong brands are likely to be able leverage those brands on the Netpage Network. They have an incentive to do so quickly to prevent newcomers from filling the vacuum and capturing the attendant advertising revenue.

News publishers also have an incentive to migrate to the Netpage Network because it allows them to offer the more fine-grained targeting that advertisers are increasingly demanding, and which they are increasingly seeking elsewhere.

- News publishers who create content rather than simply aggregating other sources have a significant advantage, since they offer both unique content and an editorial voice. Users are more likely to choose a single news publication whose content and editorial orientation they find useful, than specifying to a news aggregator how to glue together a number of disparate news sources. And any sufficiently strong news publication brand is unlikely to make its content available to an aggregator, since the aggregator will be taking a proportion of advertising and e-commerce revenue.

The Netpage Network, like the Web, offers lower barriers to entry than traditional publishing media, and this naturally stimulates greater diversity. However, the geographic independence of the network, coupled with built-in mechanisms for localization of publications, allows international, national and regional news publications to more easily compete in local news markets.

The strength of a traditional local news publication lies partly in its local news content and partly in its local retail advertising and classified advertising content. Aggregation of classified advertising is already happening on the Web, and the Netpage Network will make the same thing possible for local retail advertising. Local news publications are therefore likely to be excluded from the direct capture of local advertising, and may instead transform themselves into news gatherers feeding localized editions of larger publications.

10.3.2 Advertising Aggregators

The Netpage Network promises to be the most effective advertising medium ever conceived. It combines the editorial and print quality of traditional publications with arbitrarily finely targeted advertising, and provides a direct link between advertising, product information, and purchasing.

Because personalization and localization are handled automatically by Netpage Publication Servers, an advertising aggregator can provide arbitrarily broad coverage of both geography and demographics. The subsequent disaggregation is efficient because it is automatic.

This makes it more cost-effective for publishers to deal with advertising aggregators than to directly capture advertising. Even though the advertising aggregator is taking a proportion of advertising revenue, publishers may find the change profit-neutral because of the greater efficiency of aggregation.

Because of the finer targeting supported by the Netpage Network, publishers and advertising aggregators have a larger advertising space to sell, leading to greater profits. The linking between advertising, detailed product information, and purchasing, and the corresponding measurability of consumer behavior, leads to greater profits from click-through fees and e-commerce commissions, benefiting publishers and advertising aggregators alike.

Added revenue from these fees and commissions may even allow users' costs - printer, ink, paper, and Internet access - to be fully subsidized.

10.3.3 Netpage Points

One possible subsidization mechanism involves maintaining an account of Netpage "points" for each user.

Netpage points are earned in proportion to the amount of advertising the user accepts. Points are automatically deducted to cover the running costs of the printer. The user can also redeem points for particular products and services offered on the Netpage Network, some of which may only be available in exchange for points. Points cannot be redeemed for cash. A user who has insufficient Netpage points can be billed directly for Netpage running costs.

Netpage points may also be used more widely to reward a user for making purchases via the Netpage Network.

10.3.4 Netpage Newspaper as Portal

The front page of a personalized newspaper delivered via the Netpage Network can contain a standard set of category buttons, including for example:

- Books
- Clothing
- Electronics
- Entertainment
- Finance
- Groceries
- Health

- House and Garden
- Looking Good
- Music
- Restaurants
- Special Offers
- Sports
- Travel, and
- Emergency

Advertisers whose goods or services relate to the category of a particular button can buy the right to be linked to that button. Several levels of linking are possible. An advertiser can buy the entire link, so that the button leads directly to the advertiser's page. Or the advertiser can join a directory of advertisers who are linked to the button, so that the advertiser's page is linked to a (usually small) directory which is linked to the button. The publisher chooses how to deploy each button category and each button instance. The publisher can segment the button space of their subscriber base in arbitrary ways, for example by demographics, location, or simply proportionally to satisfy advertisers who wish to purchase a subset of the space.

Certain category buttons, such as "Emergency", can link to non-advertising information specific to the user's location, and are provided as a service by the publisher.

Web portals typically offer similar category buttons to advertisers.

REFERENCES

11 Silverbrook References

- [1] Silverbrook Research, *Artcam*, 1997
- [2] Silverbrook Research, *Memjet*, 1998
- [3] Silverbrook Research, *Netpage Applications*, v0.1, 1999

12 Other References

- [4] Adobe Systems Inc., *Portable Document Format Reference Manual*, version 1.2, November 27, 1996
- [5] Bank of Montreal Economics Department, "Newsprint Market Update", Forest Products, May 1998
- [6] Bender, W., "Read all about it in the Daily You", Communicating Business, Forward Publishing, London, Winter 1994/95 (<http://nif.www.media.mit.edu/papers/forward.html>)
- [7] Brand, S., "Personal Newspaper", in *The Media Lab*, Penguin Books, 1988, pp.36-39
- [8] Bray, T., "News Wire Services Heading for XML", 12 August 1998, <http://xml.com/xml/pub/98/08/nitf.html>
- [9] Cable Datacom News, "Cable Modem Customer Count To Top 500,000 At Year's End", December 1998, <http://www.cabledatacomnews.com/dec98-1.htm>
- [10] Cable Foreman Xtra, "WDM boosts fiber-optic capacity", May 1998, <http://www.broadband-guide.com/cbl4man/eng/eng5981.html>
- [11] Chesnais, P.R., M.J. Mucklo and J.A. Sheena, "The Fishwrap Personalized News System", *IEEE Second International Workshop on Community Networking*, Integrating Multimedia Services to the Home, 1995
- [12] Cisco Systems, "Cisco 605 Personal PCI ADSL Modem", <http://www.cisco.com/warp/public/728/605>
- [13] Cisco Systems, "Cisco 6100 Series advanced DSL Access Multiplexer", <http://www.cisco.com/warp/public/728/6100>
- [14] Clark, T., "Visa, Mastercard try to revive SET", *CNET NEWS.COM*, 12 May 1999, <http://www.news.com/News/Item/0,4,0-36435,00.html@st.ne.180.head>
- [15] Click-Through, <http://www.clickthrough.com>
- [16] CNN Interactive, "With 'electronic ink', a printed page can be a work in progress", <http://cnn.com/TECH/science/9812/23/electronic.ink.ap/index.html>
- [17] Cook, P.S., D. Gomery, and L.W. Lichty (eds.), *The Future of News - Television, Newspapers, Wire Services, Newsmagazines*, The Woodrow Wilson Center Press, 1992
- [18] Crothers, B., "IBM eyes new way to sell PCs", *CNET NEWS.COM*, 14 May 1999, <http://www.news.com/News/Item/0,4,0-36508,00.html@st.ne.ni.lh>
- [19] Daniel, R., and Mealling, M., "Resolution of Uniform Resource Identifiers using the Domain Name System", *Internet Engineering Task Force RFC2168*, June 1997
- [20] Doane, W.J., "A Comparison of Display Technologies for E-Books", <http://www.jmc.kent.edu/futureprint/1998fall/doane.htm>
- [21] E Ink, <http://www.eink.com>
- [22] Ericsson, *3G Cellular Technology*, <http://www.ericsson.se/wireless/products/modsys/3rdgen/subpages/3gcell>
- [23] E&P Interactive, "Online News Reading Soars", *E&P Interactive*, 12 June 1998, <http://www.mediainfo.com/ephome/news/newshtm/stories/061298n4.htm>
- [24] E&P Interactive, "Audience Rather than Circulation", *E&P Interactive*, 24 July 1998, <http://www.mediainfo.com/ephome/news/newshtm/stories/072498n4.htm>
- [25] E&P Interactive, "Newspapers Near Death, Gates Says", *E&P Interactive*, 5 February 1999, <http://www.mediainfo.com/ephome/news/newshtm/stories/020599n1.htm>
- [26] Fitzgerald, M., "Newspapers Re-think Microzoning", *E&P Interactive*, 14 January 1999, <http://www.mediainfo.com/ephome/news/newshtm/stories/011499n3.htm>
- [27] Gannet Co. Inc., *Consolidated Statements of Income 1998*, <http://www.gannet.com/annual/ar98/financials/income.htm>
- [28] Handley, K., "All the World's a Page", http://www.ingersoll-rand.com/compair/oct_nov_97/octnovco/pag_1.htm
- [29] Hughes Network Systems, "Hughes Spaceway - Wireless Broadband on Demand", <http://www.hns.com/spaceway/spaceway.htm>
- [30] INRIA, "Light-weight Reliable Multicast Protocol", <http://monet.inria.fr/lrmp/lrmp.html>
- [31] INRIA, "WebCanal White Paper - Global Information Broadcast", <http://monet.inria.fr/white/index.html>

- [32] IP Multicast Initiative (IPMI), *Higher Level Protocols used with IP Multicast*, 1997, [http://www.ipmulticast.com/...](http://www.ipmulticast.com/)
- [33] IPTC - NAA, *News Industry Text Format*, version 2.0b1, June 1998
- [34] ISO/IEC JTC1/SC29/WG1 (Joint Photographic Experts Group), *Call for contributions for JPEG 2000: Image Coding System*, <http://www.jpeg.org/public/wg1n505.pdf>
- [35] ISO/IEC JTC1/SC29/WG1 (Joint Photographic Experts Group), *JPEG 2000 requirements and profiles version 4.0*, <http://www.jpeg.org/public/wg1n1105.pdf>
- [36] Kent Displays, <http://www.kentdisplays.com>
- [37] Krause, J.K., "Two More 'Free' PC Firms Pin Hopes on E-commerce", *The Industry Standard*, 1 April 1999, http://www.thestandard.net/articles/article_print/0,1454,4073,00.html
- [38] Lambert, P., "Multicast Debuts on Cable Modems", *ZDNN*, 5 August 1997, <http://www5.zdnet.com/zdnn/content/inwk/0426/inwk0006.html>
- [39] Liebeskind, K., "Newspapers Continue to Report Higher Revenues", *E&P Interactive*, 8 January 1998, <http://www.mediainfo.com/ephome/news/newshtm/stories/010899n1.htm>
- [40] Loberg, C.J., "Testing the Future of All-Optical Networks", *Lightwave Xtra*, March 1998, <http://www.broadband-guide.com/lw/reports/report3981.html>
- [41] McCanne, S., "Overview of Reliable Multicast Protocols", *Matters Mbone Workshop*, SIGCOMM '96, August 1996, http://gaia.cs.umass.edu/sigcomm_mcast/talk1.html
- [42] McCanne, S., "SRM: Work in Progress", *Matters Mbone Workshop*, SIGCOMM '96, August 1996, http://gaia.cs.umass.edu/sigcomm_mcast/talk2.html
- [43] MasterCard/Visa, *SET Secure Electronic Transaction Specification, Book 1: Business Description*, Version 1.0, 31 May 1997, http://www.setco.org/set_specifications.html
- [44] Medin, M., "Data Over Cable and the @Home Network", <http://www.academ.com/nanog/feb1996/data.over.cable.html>
- [45] MIT Media Lab, *News in the future*, <http://nif.www.media.mit.edu/nif.html>
- [46] MIT Media Lab, *Fishwrap Documentation Center*, <http://fishwrap-docs.www.media.mit.edu/docs>
- [47] MIT Media Lab, *Micromedia - Electronic Paper*, <http://www.media.mit.edu/micromedia/elecpaper.html>
- [48] My Yahoo, <http://my.yahoo.com>
- [49] Negroponte, N., "Less is More", in *Being digital*, Vintage Books, 1996, pp.149-159
- [50] New Media Communication, *CyberStream Fast Internet & Data Broadcasting System via Satellite*, <http://www.nmc.com>
- [51] New York Times Co., *Consolidated Statements of Income 1998*, http://www.nytc.com/pdf/consol1_1998.pdf
- [52] Neuwirth, R., "Inserts Outpace ROP", *E&P Interactive*, 24 September 1998, <http://www.mediainfo.com/ephome/news/newshtm/stories/092598n1.htm>
- [53] Newspaper Association of America (NAA), *Facts About Newspapers 1998*, <http://www.naa.org/info/facts>
- [54] Noack, D., "Online News Audience Becomes More Mainstream", *E&P Interactive*, 15 January 1999, <http://www.mediainfo.com/ephome/news/newshtm/stories/011599n2.htm>
- [55] Schwartz, E.I., "Advertising Webonomics 101", *Wired* magazine, February 1996
- [56] Shardanand, U. and Maes, P., "Social Information Filtering: Algorithms for Automating 'Word of Mouth'", *Proceedings of ACM Computer-Human Interaction (CHI) 1995*
- [57] Steinberg, D., "Broadening the Bandwidth to the Home in USA and Canada", 15 May 1997, <http://www.specialty.com/hiband/mi-int.html>
- [58] Sterne, J., *Advertising on the Web*, Que 1997
- [59] Stevens, W.R., *TCP/IP Illustrated, Volume 1 - The Protocols*, Addison-Wesley 1994
- [60] Stone, M., "Newspapers Must Change, Fast - API Cites Serious Threat of Online Classifieds", *E&P Interactive*, 24 July 1998, <http://www.mediainfo.com/ephome/news/newshtm/stories/072498n1.htm>
- [61] Sun Microsystems, <http://www.sun.com>
- [62] Tappert, C.C., Suen, C.Y., Wakahara, T., "The State of the Art in On-Line Handwriting Recognition", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol.12, No.8, August 1990

- [63] Teledesic, *Teledesic Overview*, <http://www.teledesic.com/overview.html>
- [64] The New York Times on the Web, <http://www.nytimes.com>
- [65] United States Postal Service, *Postal Facts May 1998*,
<http://www.usps.gov/history/pfact98.htm>
- [66] Wallace, G.K., "The JPEG Still Picture Compression Standard", *Communications of the ACM*, Vol.34, No.4, April 1991, pp.30-44
- [67] Weinberger, D., "The Daily Me? No, the Daily Us", *Wired* magazine, April 1995
- [68] World Wide Web Consortium, *Extensible Markup Language (XML) 1.0*,
W3C Recommendation, 10 February 1998, <http://www.w3.org/TR/1998/REC-xml-19980210>
- [69] World Wide Web Consortium, *Extensible Stylesheet Language (XSL) Specification*,
W3C Working Draft, 21 April 1999, <http://www.w3.org/TR/WD-XS/>
- [70] World Wide Web Consortium, *HTML 4.0 Specification*,
W3C Recommendation, 24 April 1998, http://www.w3.org/TR/REL_html40/
- [71] World Wide Web Consortium, *Scalable Vector Graphics (SVG) Specification*,
W3C Working Draft, 12 April 1999, <http://www.w3.org/TR/WD-SVG/>
- [72] Zollman, P.M., "1999 Could make or Break News Web Sites", *E&P Interactive*, 16 December 1998,
<http://www.mediainfo.com/ephome/news/newshtm/stories/121698n1.htm>

Netpage Printer

Design Description

draft version 0.3, 7 September 1999



Silverbrook Research Pty Ltd
393 Darling Street, Balmain
NSW 2041 Australia
Phone: +61 2 9818 6633
Fax: +61 2 9818 6711
Email: info@silverbrook.com.au

Confidential

Contents

PRINTER PRODUCTS.....	1
1 Overview	2
2 Wallprinter	3
2.1 Wallprinter.....	3
2.2 Wallprinter Pro	8
2.3 Wallprinter Pro R	11
3 Tableprinter.....	14
3.1 Tableprinter Pro	14
3.2 Tableprinter Pro R.....	15
4 Deskprinter	16
4.1 Deskprinter Pro R	16
5 Travelprinter	19
5.1 Travelprinter R	19
6 Microprinter	23
6.1 Microprinter.....	23
PRINTER CONTROLLER.....	25
7 Memjet-Based Printing	26
8 Document Data Flow.....	27
9 Printer Controller Architecture	29
9.1 Detailed Document Data Flow	30
9.2 Print Engine Architecture	31
10 Pen Controller Architecture	32
11 Print Engine Page Format	33
11.1 Page Structure	33
11.2 Compressed Page Format.....	34
12 Print Engine	39
12.1 Halftoner/Compositor.....	39
12.2 Printhead Interface	42
12.3 Position Tag Encoder.....	43
13 Print Engine Driver	48
13.1 Graphics and Imaging Model.....	48
13.2 Two-Layer Page Buffer	48
13.3 Compositing Model.....	50
13.4 Page Compression	51
MEMJET PRINTHEAD.....	52
14 Memjet Printhead	53
14.1 The Structure of a Memjet Segment.....	53
14.2 Making Memjet Printheads out of Segments	63
15 Memjet Printhead Interface	67
15.1 Block Diagram of Printhead Interface	68

15.2	LineSyncGen Unit.....	68
15.3	Memjet Interface.....	69
15.4	Line Loader/Format Unit.....	76
15.5	Controlling a Print.....	81
REFERENCES		82
16	Silverbrook References.....	83
17	Other References	84



PRINTER PRODUCTS



1 Overview

Netpage Printers are intended for use in domestic, commercial, corporate and hospitality environments. They all utilize pagewidth Memjet printheads [4], and are based on a simple and reliable straight paper path, passing through a Memjet transfer roller printhead mechanism. In most cases the printed page is glued along one edge and adhered to the previous page to form a final bound document that can be 1 page or 500 pages thick. They all interact with the wireless Netpage Pen.

Netpage Printers come in various forms: wall-mounting, tabletop, portable, and pocket versions.

2 Wallprinter

2.1 WALLPRINTER

A low-cost, wall-mounted, base model with a duplex 8½" Memjet printhead array that accepts a full ream of US Letter paper in a vertical format as shown in Figures 1, 3 and 4. Paper is placed into a hinged top tray down onto a sprung platen and registered under edge guides before being closed. Figure 2 shows the access to the paper and ink cartridge.

A replaceable cartridge containing cyan, magenta, yellow, and infrared inks and glue is also accessible when the tray is open. It connects via a series of self-sealing connectors to hoses that transmit ink and glue to their separate locations. The cartridge consists of a thin wall drawn aluminum casing that accommodates four ink bladders and a single glue bladder into an injection molded connector base. This is a fully recyclable product with a capacity for printing and gluing 3000 pages (1500 sheets). It is protected from forgeries by use of an authentication chip [1,2].

When closed, a release mechanism allows the platen to push the paper against the pick-up roller assembly, where it is fed directly into the duplex Memjet printhead assembly. From there, the sheet passes a momentary action glue wheel with powered spike wheels, where it has glue applied to the vertical edge as it passes through. The glue wheel is capped when not in use and is operated by a powered camshaft.

The printed sheet is fed down to a binding platen that operates with a closed steel wire loop system of pulleys, runners and a powered axle. This provides the necessary speed to push the sheet forward onto the rear of a previous sheet, glue/bind it and return to the home position to accept the next printed sheet in less than 2 seconds. A motorized paper tapper assembly aligns the sheets in a simultaneous operation.

When a document is bound and finished, a powered exit hatch with a tamper sensor opens. Plastic foils work together with the hatch to direct the finished document to the back of the collection tray and feed any further documents into the tray without hitting existing ones. The collection tray is molded in clear polycarbonate and pulls out of its socket under a certain loading. Access for removing documents is provided on three sides.

The printer has a main PCB that accommodates all major circuitry components including external data jacks. A flex PCB runs from the main PCB to the paper tray and has three different color LEDs and a push button. The LEDs indicate "on", "ink out", "paper out", and "error". The push button is a "help" interface that prints out a simple instruction sheet and a compact features directory for the user. The unit is powered by an internal 110V/220V power supply that is connected before wall-mounting.




Figure 1. Wallprinter




Figure 2. Wallprinter paper and ink cartridge access

The printer has several metal hangers on the rear, which locate into keyhole slots in a metal back plate that is securely fastened to a wall. When the printer has been connected, it is hung onto the back plate and fixed with a locking screw found under the paper tray.

The Netpage Printers are fully customizable in finishes and color as the front moldings clip on to a core chassis and are easily removable.




Figure 3. Wallprinter front elevation




Figure 4. Wallprinter side elevation

2.2 WALLPRINTER PRO

This printer is similar to Wallprinter in most respects, except that it has a duplex 11" print-head assembly, which prints on US Letter paper in a landscape format (see Figures 5, 6 and 7). This means a faster print time and binding time for each page, making for faster overall document delivery.

Another difference is the location of the ink cartridge, which resides above the paper tray rather than down the side. Each page is glued along the horizontal edge by a full-length glue sponge, which is capped when not in use. Operation, printing, and document handling are identical to Wallprinter.

Wallprinter Pro is fully customizable in finishes and color.




Figure 5. Wallprinter Pro




Figure 6. Wallprinter Pro front elevation




Figure 7. Wallprinter Pro side elevation

2.3 WALLPRINTER PRO R

This printer shares the same printing and binding configuration of US Letter landscape format as Wallprinter Pro. The main difference is the media delivery, which is in the form of a large print cartridge cartridge (see Figures 8, 9 and 10). This cartridge accommodates C, M, Y, and infrared inks and glue as well as a 1000 sheet capacity roll of paper. The cartridge can be recharged at nominated outlets when required and it is protected from forgeries by an authentication chip [1,2].

The printer has integral structural metalwork to support its weight and a ball bearing track for easy loading and removal.

Wallprinter Pro R is fully customizable in finishes and color.




Figure 8. Wallprinter Pro R, with print cartridge extracted




Figure 9. Wallprinter Pro R front elevation




Figure 10. Wallprinter Pro R side elevation

3 Tableprinter

3.1 TABLEPRINTER PRO

This printer is a tabletop version of the Wallprinter Pro. Essentially, it is the same printer unit with a base plinth that adds extra functionality, such as USB, parallel port and a power socket (see Figure 11).

Tableprinter Pro is fully customizable in finishes and color.




Figure 11. Tableprinter Pro

3.2 TABLEPRINTER PRO R

This printer is a tabletop version of the Wallprinter Pro R. Essentially, it is the same printer unit with a base plinth that adds extra functionality, such as USB, parallel port and a power socket (see Figure 12).

Tableprinter Pro R is fully customizable in finishes and color.




Figure 12. Tableprinter Pro R

4 Deskprinter

4.1 DESKPRINTER PRO R

The Deskprinter Pro R is a vertical format printer based around the large format print cartridge (see Figures 13, 14, 15 and 16). This cartridge accommodates C, M, Y, Infrared inks and glue, as well as a 1000 page roll of paper. The print cartridge is lowered into the top of the unit through a latching hatch. When printing, the paper is fed through the printhead assembly where it is duplex printed and then ejected past a cutter which crops it into sheets. As the sheet passes the cutter, a glue wheel assembly glues along the long edge and powered spike rollers propel the sheet out into the binding area.

The binding area consists of a binding platen which operates with a closed steel wire loop system of pulleys, runners and a powered axle. This provides the necessary speed to push the glued sheet forward onto the rear of a previous sheet, glue/bind it and return to the home position to accept the next printed sheet in less than 2 seconds. A motorized paper taper assembly aligns the sheets in a simultaneous operation.

When the document has been bound, a series of metal fingers at the top and bottom of the document rotate out of the way and the document is pushed into the out tray area. A powered hatch opens and a motorized ejector assembly pushes the document through the exit slot of the unit. The hatch then closes onto the document ready for removal. Subsequent documents are stacked up in a similar fashion. The user interface is identical to the other printers in the range. A large PCB contains all the processing and external interface components and connects to an internal 110V/220V power supply unit.

Deskprinter Pro R is fully customizable in finishes and color.




Figure 13. Deskprinter Pro R




Figure 14. Deskprinter Pro R, with print cartridge extracted




Figure 15. Deskprinter Pro R plan




Figure 16. Deskprinter Pro R elevation

5 Travelprinter

5.1 TRAVELPRINTER R

The Travelprinter R is a small, lightweight, versatile and completely portable Netpage Printer (see Figures 17, 19, 20 and 21). It has in-built mobile network communication hardware and software, allowing it to download documents anywhere. Travelprinter also has communications ports for computer interface printing when required.

The printer consists of a front and rear molding with a chassis to accommodate the major components including a lithium battery and an 8½" duplex Memjet printhead assembly.

A compact print cartridge with C, M, Y and infrared inks and paper is used in the printer, providing 50 US Letter pages or 100 A5 pages. It is protected from forgeries by an authentication chip.




Figure 17. Travelprinter R

A motorized guillotine assembly cuts the media between the cartridge and the printhead and motorized spike wheels eject the finished print out of the unit. A flex PCB runs from the main board to a segment LCD and two push buttons. The LCD shows signal strength, any errors, battery and number of pages left in the cartridge. The buttons allow the printer to either connect to the Netpage Network or to act as a stand-alone printer. A USB interface is provided on the side of the printer along with DC 3V input.




Figure 18. Travelprinter R main and flex PCB, with buttons and segment LCD




Figure 19. Travelprinter R, with print cartridge extracted




Figure 20. Travelprinter R detail, showing duplexed imaging units and print cartridge




Figure 21. Travelprinter plan and elevations

6 Microprinter

6.1 MICROPRINTER

The Microprinter is a small, versatile, pocket-size printer/camera (see Figures 22 and 23). It has built-in mobile network communication hardware and software, allowing it to link to the Netpage Network and fetch documents from anywhere. In addition, the product is ergonomically configured and styled as a fully functional digital camera with the standard Classic, HDTV and Panoramic print formats for photography.

The Microprinter accommodates a 4" (100mm) page-width print cartridge with C, M, Y and infrared ink plus 5.4 meters of paper. This means the user can single-sided print 41 Netpages, 36 Classic, 30 Horizontal or 18 Panoramic prints or any combination thereof.




Figure 22: Microprinter rear view with printing in progress (left), and front view with print cartridge extracted (right)

The Microprinter is a fixed-lens camera with auto focus and digital photo enhancement capabilities, which allows the user to take sophisticated photos very simply and easily. A print button allows the user to print and reprint a photo or Netpage.

The Microprinter consists of two main front and rear moldings and a hatch. The main moldings accommodate the lens assembly, viewfinder optics and a rigid PCB. A flex PCB runs from the main board to the imaging chip, an LCD, the printhead assembly, the print-

head capping mechanism, a cutter assembly, the roll motor drive, various control switches and a CR2 battery.

The Microprinter is fully customizable in finishes and color.




Figure 23: Microprinter exploded view

PRINTER CONTROLLER

7 Memjet-Based Printing

A Memjet printhead produces 1600 dpi bi-level CMYK. On low-diffusion paper, each ejected drop forms an almost perfectly circular 22.5µm diameter dot. Dots are easily produced in isolation, allowing dispersed-dot dithering to be exploited to its fullest. Since the Memjet printhead is the width of the page and operates with a constant paper velocity, the four color planes are printed in perfect registration, allowing ideal dot-on-dot printing. Since there is consequently no spatial interaction between color planes, the same dither matrix is used for each color plane. Dot-on-dot printing minimizes ‘muddying’ of mid-tones caused by inter-color bleed.

A page layout may contain a mixture of images, graphics and text. Continuous-tone (contone) images and graphics are reproduced using a stochastic dispersed-dot dither. Unlike a clustered-dot (or amplitude-modulated) dither, a *dispersed-dot* (or frequency-modulated) dither reproduces high spatial frequencies (i.e. image detail) almost to the limits of the dot resolution, while simultaneously reproducing lower spatial frequencies to their full color depth, when spatially integrated by the eye. A *stochastic* dither matrix is carefully designed to be free of objectionable low-frequency patterns when tiled across the image. As such its size typically exceeds the minimum size required to support a particular number of intensity levels (e.g. 16×16×8 bits for 257 intensity levels).

Human contrast sensitivity peaks at a spatial frequency of about 3 cycles per degree of visual field and then falls off logarithmically, decreasing by a factor of 100 beyond about 40 cycles per degree and becoming immeasurable beyond 60 cycles per degree [7,8]. At a normal viewing distance of 12 inches (about 300mm), this translates roughly to 200-300 cycles per inch (cpi) on the printed page, or 400-600 samples per inch according to Nyquist’s theorem.

In practice, contone resolution above about 300 ppi is of limited utility outside special applications such as medical imaging. Offset printing of magazines, for example, uses contone resolutions in the range 150 to 300 ppi. Higher resolutions contribute slightly to color error through the dither.

Black text and graphics are reproduced directly using bi-level black dots, and are therefore not antialiased (i.e. low-pass filtered) before being printed. Text is therefore *supersampled* beyond the perceptual limits discussed above, to produce smoother edges when spatially integrated by the eye. Text resolution up to about 1200 dpi continues to contribute to perceived text sharpness (assuming low-diffusion paper, of course).

The Netpage Printer uses a contone resolution of 267 ppi (i.e. 1600 dpi / 6), and a black text and graphics resolution of 800 dpi.

8 Document Data Flow

Because of the page-width nature of the Memjet printhead, each page must be printed at a constant speed to avoid creating visible artifacts. This means that the printing speed can't be varied to match the input data rate. Document rasterization and document printing are therefore decoupled to ensure the printhead has a constant supply of data. A page is never printed until it is fully rasterized. This is achieved by storing a compressed version of each rasterized page image in memory.

This decoupling also allows the RIP to run ahead of the printer when rasterizing simple pages, buying time to rasterize more complex pages.

Because contone color images are reproduced by stochastic dithering, but black text and line graphics are reproduced directly using dots, the compressed page image format contains a separate foreground bi-level black layer and background contone color layer. The black layer is composited over the contone layer after the contone layer is dithered.

Figure 24 shows the flow of a Netpage Printer document from network to printed page.

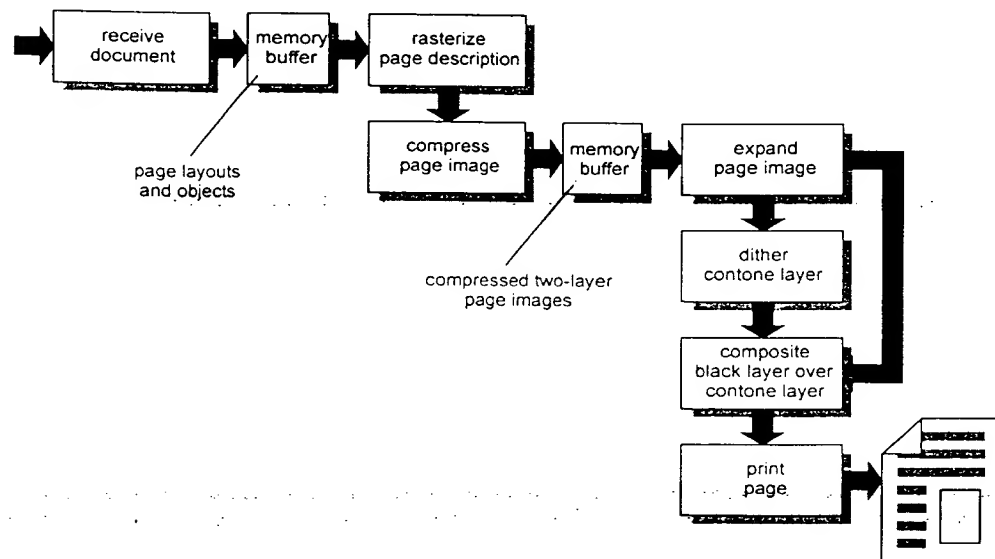


Figure 24. Netpage Printer document data flow

At 267 ppi, a Letter page of contone RGB or CMY data has a size of 19MB. Using lossy contone compression algorithms such as JPEG [9], contone images compress with a ratio up to 10:1 without noticeable loss of quality, giving a compressed page size of 1.9MB.

At 800 dpi, a Letter page of bi-level data has a size of 7MB. Coherent data such as text compresses very well. Using lossless bi-level compression algorithms such as Group 4 Facsimile [6], ten-point text compresses with a ratio of about 10:1 (as discussed in Section 11.2.1.2), giving a compressed page size of 0.8MB.

Once dithered, a page of CMY contone image data consists of 86MB of bi-level data. Using lossless bi-level compression algorithms on this data is pointless precisely because the optimal dither is stochastic - i.e. since it introduces hard-to-compress disorder.

The two-layer compressed page image format therefore exploits the relative strengths of lossy JPEG contone image compression and lossless bi-level text compression. The format is compact enough to be storage-efficient, and simple enough to allow straightforward real-time expansion during printing.

Since text and images normally don't overlap, the normal worst-case page image size is 1.9MB (i.e. image only), while the normal best-case page image size is 0.8MB (i.e. text only). The absolute worst-case page image size is 2.7MB (i.e. text over image). Assuming a quarter of an average page contains images, the average page image size is 1.1MB.

9 Printer Controller Architecture

The Netpage Printer controller consists of a controlling processor, a factory-selected network interface, a radio transceiver, dual raster image processor (RIP) DSPs, duplexed print engines, flash memory, and 64MB of DRAM, as illustrated in Figure 25.

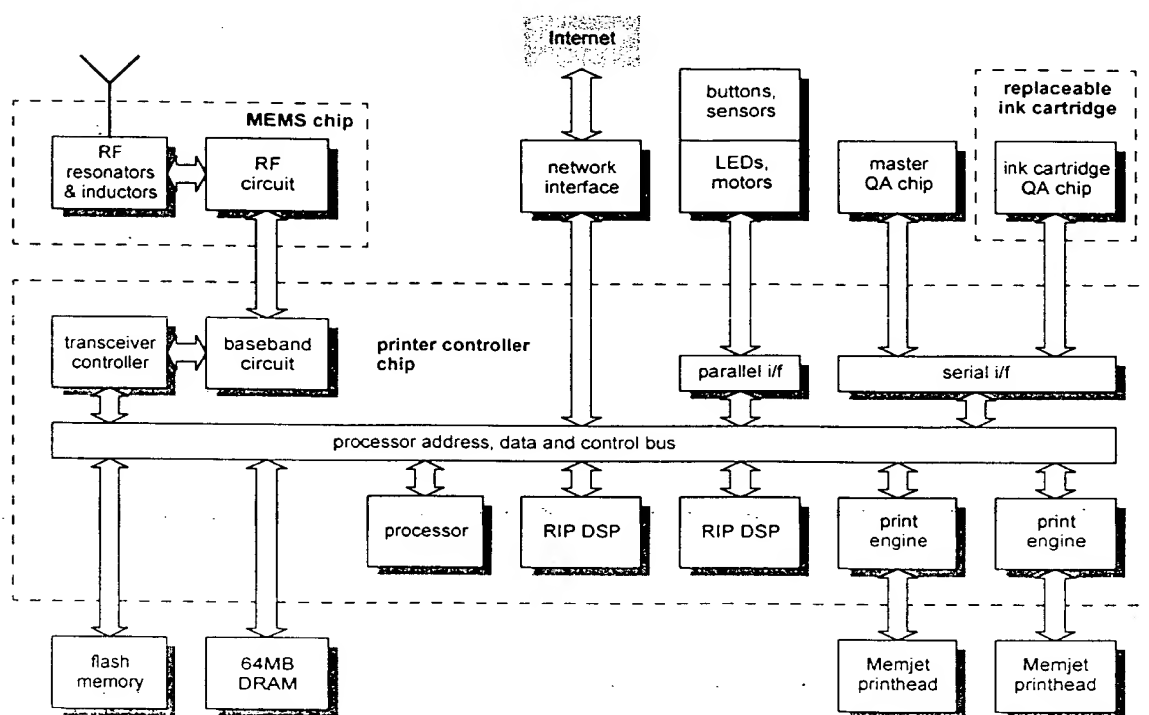


Figure 25. Basic printer controller architecture

The controlling processor handles communication with the Internet and with local wireless pens, controls the user interface (buttons and LEDs), controls the paper transport, handles ink cartridge authentication and ink monitoring, and feeds and synchronizes the RIP DSPs and print engines. It consists of a medium-performance general-purpose microprocessor.

The RIP DSPs rasterize and compress page descriptions to the Netpage Printer's compressed page format. Each print engine expands, dithers and prints page images to its associated Memjet printhead in real time (i.e. at 30 or 45 pages per minute). The duplexed print engines print both sides of the page simultaneously.

The printer controller's flash memory holds the software for both the processor and the DSPs, as well as configuration data. This is copied to main memory at boot time.

The processor, DSPs, print engines and digital transceiver components are integrated in a single ASIC. The MEMS and analog RF components are integrated in a separate MEMS chip, which is also used in the wireless pen. Additional pen-specific components in the

MEMS chip are not used in the printer controller. The Internet network interface module is separate, since Netpage Printers allow the network connection to be factory-selected. Flash memory and the 2x256Mbit (64MB) DRAM is also off-chip.

Various Internet network interface modules can be supported. Possibilities include a POTS modem, a Hybrid Fiber-Coax (HFC) cable modem, an ISDN modem, a DSL modem, a satellite transceiver, a current or next-generation cellular telephone transceiver, a wireless local loop (WLL) transceiver, etc. A Internet connection may already be available on the local network, in which case only a local network connection may be required.

The printer controller optionally includes a local network connection, to allow the printer to be used from a directly-connected workstation or over a local-area network. Possibilities include 10Base-T and 100Base-T Ethernet, USB and USB 2.0, IEEE 1394 (Firewire), and various emerging home networking standards.

The radio transceiver communicates in the unlicensed 900MHz band normally used by cordless telephones, and uses frequency hopping and collision detection to provide interference-free communication.

9.1 DETAILED DOCUMENT DATA FLOW

The main processor receives and verifies the document's page layouts and page objects by Internet pointcast and multicast. It then runs the appropriate RIP software on the DSPs.

The DSPs rasterize each page description and compress the rasterized page image. The main processor stores each compressed page image in memory. The simplest way to load-balance multiple DSPs is to let each DSP rasterize a separate page. The DSPs can always be kept busy since an arbitrary number of rasterized pages can, in general, be stored in memory. This strategy can lead to poor DSP utilization, however, when rasterizing short documents.

Watermark regions in the page description are rasterized to a contone-resolution bi-level bitmap which is losslessly compressed to negligible size and which forms part of the compressed page image.

The infrared (IR) layer of the printed page contains encoded position tags at a density of about 25 per inch. Each tag encodes the page id, tag position, and control bits. Active areas and pressure-sensitive areas in the page description are rasterized to tag-resolution bi-level bitmaps which do not require compression and which form part of the compressed page image.

The main processor passes back-to-back page images to the duplexed print engines. Each print engine stores the compressed page image into its local memory, and starts the page expansion and printing pipeline. Page expansion and printing is pipelined because it is impractical to store a 114MB bi-level CMY+IR page image in memory.

The first stage of the pipeline expands the JPEG-compressed contone CMY layer, expands the Group 4 Fax-compressed bi-level watermark map, and expands the Group 4 Fax-compressed bi-level black layer, all in parallel. The second stage dithers the contone CMY layer using the dither matrix selected by the watermark map, and composites the bi-level black layer over the resulting bi-level CMY layer. Since there is no black ink used in the Netpage Printer, the black layer is composited with each of C, M and Y. In parallel with

this, the tag encoder generates and encodes the bi-level IR tag data. The last stage prints the bi-level CMY+IR data through the Memjet printhead via the printhead interface.

9.2 PRINT ENGINE ARCHITECTURE

The print engine's page expansion and printing pipeline consists of a standard JPEG decoder, a standard Group 4 Fax decoder, a custom halftoner/compositor unit, a custom position tag encoder, and a custom interface to the Memjet printhead. These are described in detail in Section 12.

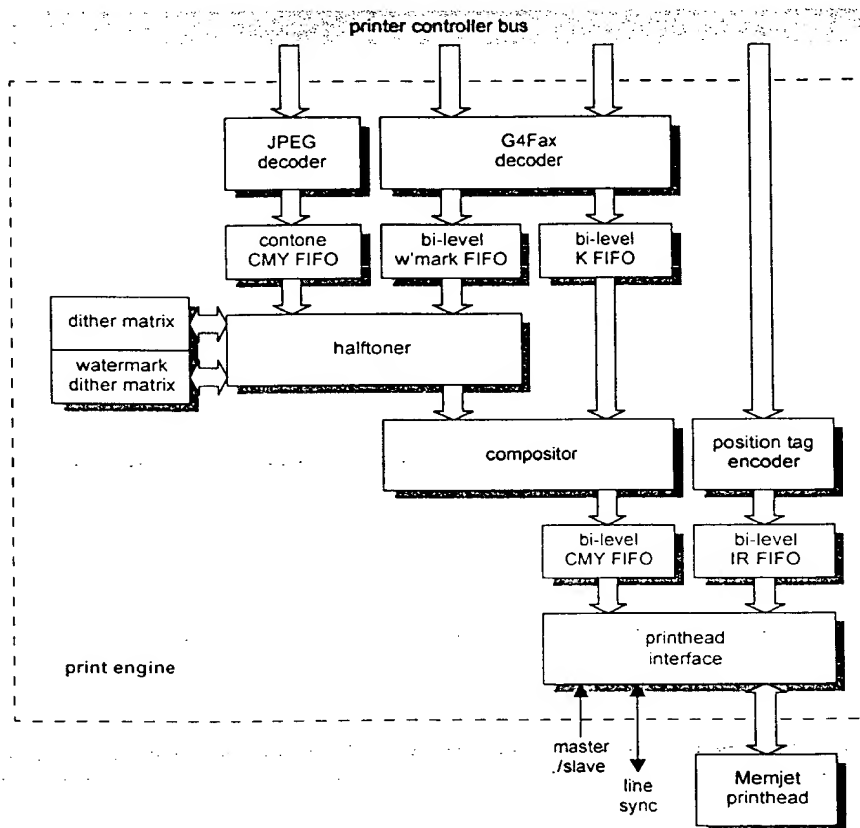


Figure 26. Print engine architecture

When several print engines are used in unison, such as in a duplexed configuration, they are synchronized via a shared line sync signal. Only one print engine, selected via the external master/slave pin, generates the line sync signal onto the shared line.

In the 8½" versions of the Netpage Printer, the two print engines each prints 30 Letter pages per minute along the long dimension of the page (11"), giving a line rate of 8.8 kHz at 1600 dpi. In the 11" Pro versions of the Netpage Printer, the two print engines each prints 45 Letter pages per minute along the short dimension of the page (8½"), giving a line rate of 10.2 kHz. These line rates are well within the operating frequency of the Memjet printhead, which in the current design exceeds 30 kHz.

10 Pen Controller Architecture

The Netpage Pen controller consists of a controlling processor, a radio transceiver, a tilt sensor, a nib pressure sensor, a IR image sensor, flash memory, and 512KB of DRAM, as illustrated in Figure 27.

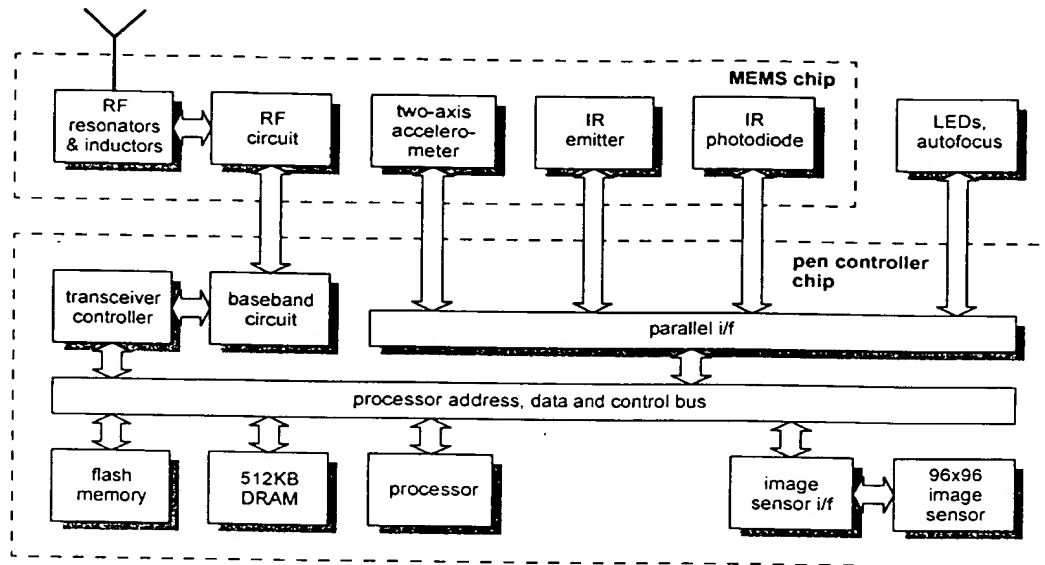


Figure 27. Pen controller architecture

The controlling processor captures and decodes IR position tags from the page via the image sensor, monitors the tilt and pressure sensors, controls the autofocus voice-coil, controls the user interface (tricolor LED), and handles wireless communication with the local Netpage Printer. It consists of a medium-performance (~25MHz RISC) general-purpose microprocessor.

Two-axis tilt sensing is provided by a two-axis accelerometer. Nib pressure sensing is provided by an IR emitter and photodiode pair in conjunction with a reflector coupled with the sprung nib.

The processor, digital transceiver components, 96×96 image sensor, flash memory and 512KB DRAM are integrated in a single ASIC. The MEMS and analog RF components, accelerometers, and the IR emitter/photodiode are integrated in a single MEMS chip, also used in the Netpage Printer.

The radio transceiver communicates in the unlicensed 900MHz band normally used by cordless telephones, and uses frequency hopping and collision detection to provide interference-free communication.

11 Print Engine Page Format

This section describes the format of compressed pages expected by the print engines.

The raster image processors (RIPs) generate pages in this format. The compressed format and the print engines are designed to allow real-time page expansion during printing, to ensure that printing is never interrupted in the middle of a page due to data underrun.

The Netpage Printer reproduces black at full dot resolution (1600 dpi), but reproduces contone color at a somewhat lower resolution using halftoning. The page description is therefore divided into a black layer and a contone layer. The black layer is defined to composite *over* the contone layer.

The black layer consists of a bitmap containing a 1-bit *opacity* for each pixel. This black layer *matte* has a resolution which is an integer factor of the printer's dot resolution. The highest supported resolution is 1600 dpi, i.e. the printer's full dot resolution.

The contone layer consists of a bitmap containing a 24-bit CMY *color* for each pixel. This contone image has a resolution which is an integer factor of the printer's dot resolution. The highest supported resolution is 267 ppi, i.e. one-sixth the printer's dot resolution.

The contone resolution is also typically an integer factor of the black resolution, to simplify calculations in the RIPs. This is not a requirement, however.

The black layer and the contone layer are both in compressed form for efficient storage in the printer's internal memory.

11.1 PAGE STRUCTURE

The Netpage Printer prints with full edge bleed using an 8½" printhead. It imposes no margins and so has a printable page area which corresponds to the size of its paper. The target page size is constrained by the printable page area, less the explicit (target) left and top margins specified in the page description. These relationships are illustrated below.




Figure 28. Page structure

11.2 COMPRESSED PAGE FORMAT

Apart from being implicitly defined in relation to the printable page area, each page description is complete and self-contained. There is no data stored separately from the page description to which the page description refers.

The page description consists of a page header which describes the size and resolution of the page, followed by one or more page bands which describe the actual page content.

Table 1 shows the format of the page header.

Table 1. Page header format

field	format	description
signature	16-bit integer	Page header format signature.
version	16-bit integer	Page header format version number.
structure size	16-bit integer	Size of page header.
target resolution (dpi)	16-bit integer	Resolution of target page. This is always 1600 for the Netpage Printer.
target page width	16-bit integer	Width of target page, in dots.
target page height	16-bit integer	Height of target page, in dots.
target left margin	16-bit integer	Width of target left margin, in dots.
target top margin	16-bit integer	Height of target top margin, in dots.
black scale factor	16-bit integer	Scale factor from black resolution to target resolution (must be 2 or greater).
black page width	16-bit integer	Width of black page, in black pixels.
black page height	16-bit integer	Height of black page, in black pixels.
contone scale factor	16-bit integer	Scale factor from contone resolution to target resolution (must be 6 or greater).
contone page width	16-bit integer	Width of contone page, in contone pixels.
contone page height	16-bit integer	Height of contone page, in contone pixels.

The page header contains a signature and version which allow the print engine to identify the page header format. If the signature and/or version are missing or incompatible with the print engine, then the print engine can reject the page.

The page header defines the resolution and size of the target page. The black and contone layers are clipped to the target page if necessary. This happens whenever the black or contone scale factors are not factors of the target page width or height.

The target left and top margins define the positioning of the target page within the printable page area.

The black layer parameters define the pixel size of the black layer, and its integer scale factor to the target resolution.

The contone layer parameters define the pixel size of the contone layer, and its integer scale factor to the target resolution.

Table 2 shows the format of the page band header.

Table 2. Page band header format

field	format	description
signature	16-bit integer	Page band header format signature.
version	16-bit integer	Page band header format version number.
structure size	16-bit integer	Size of page band header.
black band height	16-bit integer	Height of black band, in black pixels.
black band data size	32-bit integer	Size of black band data, in bytes.
contone band height	16-bit integer	Height of contone band, in contone pixels.
contone band data size	32-bit integer	Size of contone band data, in bytes.
watermark map band data size	32-bit integer	Size of watermark map band data, in bytes.
tag control band height	16-bit integer	Height of position tag control band, in tags.
active area map band size	32-bit integer	Size of active area map band data, in bytes.
pressure-sensitive area map band size	32-bit integer	Size of pressure-sensitive area band data, in bytes.

The black layer parameters define the height of the black band, and the size of its compressed band data. The variable-size black data follows the page band header.

The contone layer parameters define the height of the contone band, and the size of its compressed page data, consisting of the contone color data and the associated bi-level watermark map. The variable-size contone data follows the black data. The variable-size bi-level watermark map data follows the contone data.

The tag layer parameters define the height of the position tag control band, and the size of its active area and pressure-sensitive area map band data. The variable-size active area map data follows the watermark map data. The variable-size pressure-sensitive area map data follows the active area map data.

Table 3 shows the format of the variable-size compressed band data which follows the page band header.

Table 3. Page band data format

field	format	description
black data	G4Fax bytestream	Compressed bi-level black data.
contone data	JPEG bytestream	Compressed contone CMY data.
watermark map	G4Fax bytestream	Compressed bi-level watermark map data.
active area map	bjitmap	Bi-level active area map data.
pressure-sensitive area map	bitmap	Bi-level pressure-sensitive area map data.

Each variable-size segment of band data is aligned to an 8-byte boundary.

The following sections describe the format of the compressed bi-level layers and the compressed contone layer.

11.2.1 Bi-level Data Compression

11.2.1.1 Group 3 and 4 Facsimile Compression

The Group 3 Facsimile compression algorithm [12] losslessly compresses bi-level data for transmission over slow and noisy telephone lines. The bi-level data represents scanned black text and graphics on a white background, and the algorithm is tuned for this class of images (it is explicitly not tuned, for example, for *half-toned* bi-level images). The 1D Group 3 algorithm runlength-encodes each scanline and then Huffman-encodes the resulting runlengths. Runlengths in the range 0 to 63 are coded with *terminating* codes. Runlengths in the range 64 to 2623 are coded with *make-up* codes, each representing a multiple of 64, followed by a terminating code. Runlengths exceeding 2623 are coded with multiple make-up codes followed by a terminating code. The Huffman tables are fixed, but are separately tuned for black and white runs (except for make-up codes above 1728, which are common). When possible, the 2D Group 3 algorithm encodes a scanline as a set of short edge deltas (0, ± 1 , ± 2 , ± 3) with reference to the previous scanline. The delta symbols are entropy-encoded (so that the zero delta symbol is only one bit long etc.) Edges within a 2D-encoded line which can't be delta-encoded are runlength-encoded, and are identified by a prefix. 1D- and 2D-encoded lines are marked differently. 1D-encoded lines are generated at regular intervals, whether actually required or not, to ensure that the decoder can recover from line noise with minimal image degradation. 2D Group 3 achieves compression ratios of up to 6:1 [12].

The Group 4 Facsimile algorithm [6] losslessly compresses bi-level data for transmission over *error-free* communications lines (i.e. the lines are truly error-free, or error-correction is done at a lower protocol level). The Group 4 algorithm is based on the 2D Group 3 algorithm, with the essential modification that since transmission is assumed to be error-free, 1D-encoded lines are no longer generated at regular intervals as an aid to error-recovery. Group 4 achieves compression ratios ranging from 20:1 to 60:1 for the CCITT set of test images [12].

The design goals and performance of the Group 4 compression algorithm qualify it as a compression algorithm for the bi-level layers. However, its Huffman tables are tuned to a lower scanning resolution (100-400 dpi), and it encodes runlengths exceeding 2623 awkwardly. At 800 dpi, our maximum runlength is currently 6400. Although a Group 4 decoder core might be available for use in the printer controller chip, it might not handle runlengths exceeding those normally encountered in 400 dpi facsimile applications, and so would require modification.

11.2.1.2 Edge Delta / Runlength (EDRL) Compression

Since most of the benefit of Group 4 comes from the delta-encoding, a simpler algorithm based on delta-encoding alone is likely to meet our requirements. This approach, in the form of the Edge Delta / Runlength (EDRL) encoding, is described in detail in another Silverbrook design document [3]. There are therefore two viable approaches to losslessly compressing bi-level data: Group 4 Facsimile and EDRL. Their compression performance is compared below. Where this document refers to Group 4 Facsimile (G4Fax) compression, EDRL can be substituted, purely as an implementation decision.

Magazine text is typically typeset in a typeface with serifs (such as Times) at a point size of 10. At this size a Letter page holds up to 14,000 characters, though a typical magazine page holds only about 7,000 characters. Text is seldom typeset at a point size smaller than 5. At 800 dpi, text cannot be meaningfully rendered at a point size lower than 2 using a standard typeface. Table 4 illustrates the legibility of various point sizes.

Table 4. Text at different point sizes

point size	sample text (in Times)
2	The quick brown fox jumps over the lazy dog.
3	The quick brown fox jumps over the lazy dog.
4	The quick brown fox jumps over the lazy dog.
5	The quick brown fox jumps over the lazy dog.
6	The quick brown fox jumps over the lazy dog.
7	The quick brown fox jumps over the lazy dog.
8	The quick brown fox jumps over the lazy dog.
9	The quick brown fox jumps over the lazy dog.
10	The quick brown fox jumps over the lazy dog.

Table 5 shows Group 4 and EDRL compression performance on pages of text of varying point sizes, rendered at 800 dpi. The distribution of characters on the test pages is based on English-language statistics [11].

Table 5. Group 4 and EDRL compression performance on text at 800 dpi

point size	characters per page	Group 4 compression ratio	EDRL compression ratio
2	340,000	2.3	1.7
3	170,000	3.2	2.5
4	86,000	4.7	3.8
5	59,000	5.5	4.9
6	41,000	6.5	6.1
7	28,000	7.7	7.4
8	21,000	9.1	9.0
9	17,000	10.2	10.4
10	14,000	10.9	11.3
11	12,000	11.5	12.4
12	8,900	13.5	14.8
13	8,200	13.5	15.0
14	7,000	14.6	16.6
15	5,800	16.1	18.5
20	3,400	19.8	23.9

For a point size of 9 or greater, EDRL slightly outperforms Group 4, simply because Group 4's runlength codes are tuned to 400 dpi.

These compression results bear out the observation that entropy-encoded runlengths contribute much less to compression than 2D encoding, unless the data is poorly correlated vertically, such as in the case of very small characters.

11.2.1.3 Black Layer Compression Format

The 800 dpi black layer is losslessly compressed using G4Fax (or EDRL) at a typical compression ratio exceeding 10:1.

11.2.1.4 Watermark Map Compression Format

The 267 dpi watermark layer, which matches the 267 ppi contone color layer, is losslessly compressed using G4Fax (or EDRL) at a typical compression ratio exceeding 50:1.

11.2.2 Contone Data Compression

11.2.2.1 JPEG Compression

The JPEG compression algorithm [9] lossily compresses a contone image at a specified quality level. It introduces imperceptible image degradation at compression ratios below 5:1, and negligible image degradation at compression ratios below 10:1 [13].

JPEG typically first transforms the image into a color space which separates luminance and chrominance into separate color channels. This allows the chrominance channels to be subsampled without appreciable loss because of the human visual system's relatively greater sensitivity to luminance than chrominance. After this first step, each color channel is compressed separately.

The image is divided into 8x8 pixel blocks. Each block is then transformed into the frequency domain via a discrete cosine transform (DCT). This transformation has the effect of concentrating image energy in relatively lower-frequency coefficients, which allows higher-frequency coefficients to be more crudely quantized. This quantization is the principal source of compression in JPEG. Further compression is achieved by ordering coefficients by frequency to maximize the likelihood of adjacent zero coefficients, and then runlength-encoding runs of zeroes. Finally, the runlengths and non-zero frequency coefficients are entropy coded. Decompression is the inverse process of compression.

11.2.2.2 CMY Contone Layer Compression Format

The CMY contone layer is compressed to an interleaved color JPEG bytestream. The interleaving is required for space-efficient decompression in the printer, but may restrict the decoder to two sets of Huffman tables rather than four (i.e. one per color channel) [13]. If luminance and chrominance are separated, then the luminance channels can share one set of tables, and the chrominance channels the other set.

If luminance/chrominance separation is deemed necessary, either for the purposes of table sharing or for chrominance subsampling, then CMY is converted to YCrCb and Cr and Cb are duly subsampled.

The JPEG bytestream is complete and self-contained. It contains all data required for decompression, including quantization and Huffman tables.



12 Print Engine

12.1 HALFTONER/COMPOSITOR

The halftoner/compositor unit (HCU) (Figure 30) combines the functions of halftoning the contone CMY layer to bi-level CMY, and compositing the black layer over the halftoned contone layer. It also selects between the normal and the “watermark” dither matrix on a pixel-by-pixel basis, based on the corresponding value in the watermark map.

The input to the HCU is an expanded 267 ppi CMY contone layer, an expanded 267 dpi watermark map, and an expanded 1600 dpi black layer. The output from the HCU is a set of 1600 dpi bi-level CMY image lines.

Once started, the HCU proceeds until it detects an *end-of-page* condition, or until it is explicitly stopped via its control register.

The HCU generates a page of dots of a specified width and length. The width and length must be written to the *page width* and *page length* registers prior to starting the HCU. The page width corresponds to the width of the printhead. The page length corresponds to the length of the target page.

The HCU generates target page data between specified left and right margins relative to the page width. The positions of the left and right margins must be written to the *left margin* and *right margin* registers prior to starting the HCU. The distance from the left margin to the right margin corresponds to the target page width.

The HCU consumes black and contone/watermark data according to specified black and contone page widths. These page widths must be written to the *black page width* and *contone page width* registers prior to starting the HCU. The HCU clips black and contone data to the target page width. This allows the black and contone page widths to exceed the target page width without requiring any special end-of-line logic at the input FIFO level.

The relationships between the page width, the black and contone page widths, and the margins are illustrated in Figure 29.




Figure 29. Relationships between page widths and margins

The HCU scales contone data to printer resolution both horizontally and vertically based on a specified scale factor. This scale factor must be written to the *contone scale factor* register prior to starting the HCU.

Table 6. Halftoner/compositor control and configuration registers

register	width	description
start	1	Start the HCU.
stop	1	Stop the HCU.
page width	14	Page width of printed page, in dots. This is the number of dots which have to be generated for each line.
left margin	14	Position of left margin, in dots.
right margin	14	Position of right margin, in dots.
page length	15	Page length of printed page, in dots. This is the number of lines which have to be generated for each page.
black page width	14	Page width of black layer, in dots. Used to detect the end of a black line.
contone page width	14	Page width of contone layer, in dots. Used to detect the end of a contone line.
contone scale factor	4	Scale factor used to scale contone data to bi-level resolution.

The consumer of the data produced by the HCU is the printhead interface. The printhead interface requires bi-level CMY image data in *planar* format, i.e. with the color planes separated. Further, it also requires that even and odd pixels are separated. The output stage of the HCU therefore uses 6 parallel pixel FIFOs, one each for *even cyan*, *odd cyan*, *even magenta*, *odd magenta*, *even yellow*, and *odd yellow*.

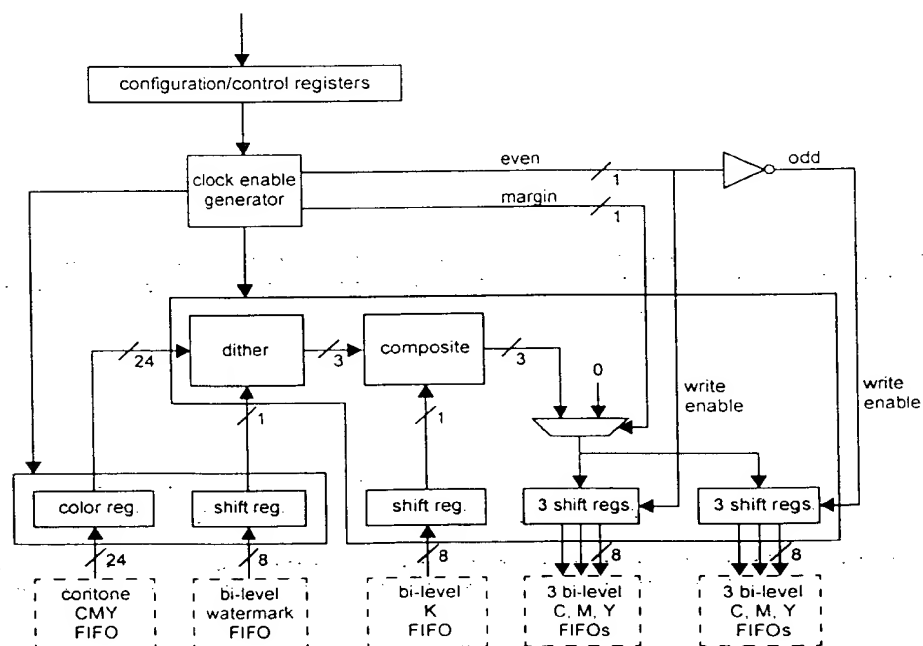


Figure 30. Halftoner/compositor unit

The input contone CMY FIFO is a full 7KB line buffer. The line is used *contone scale factor* times to effect vertical up-scaling via line replication. FIFO write address wrapping is disabled until the start of the last use of the line.

12.1.1 Dither

The dither unit converts a 24-bit contone CMY value into a 3-bit bi-level CMY value. Each color component is independently compared with the 8-bit threshold value at the current dot position in the dither cell (see Figure 31).

There are two programmable 64×64×8-bit dither cells: a normal dither cell and a “watermark” dither cell. The 1-bit watermark value associated with the current contone pixel controls the selection of the watermark or normal threshold value.

The dither cell address generator contains a pair of 6-bit row and column address registers. Both registers are initialized to zero. The row address is incremented on every dot line advance, modulo 64. The column address is incremented on every dot clock, modulo 64, and is reset on a dot line advance.

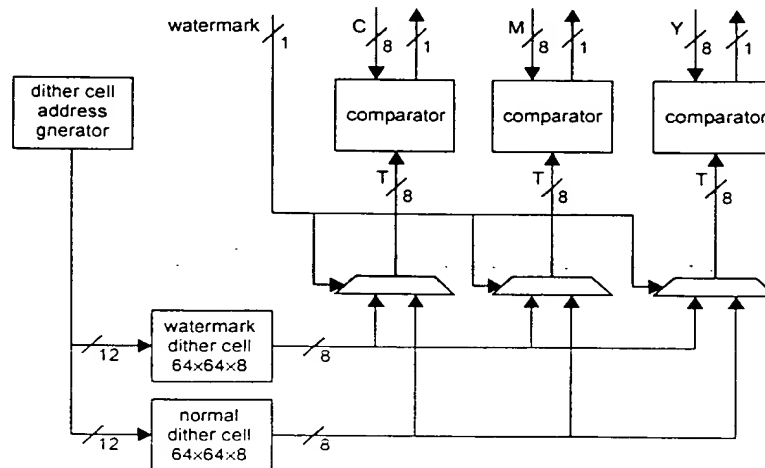


Figure 31. Dither, with watermarking

12.1.2 Composite

The composite unit composites a black layer dot over a halftoned CMY layer dot. If the black layer opacity is one, then the halftoned CMY is set to black (i.e. all ones). The black layer opacity is therefore simply ORed with each of the halftoned C, M, and Y values.

12.1.3 Clock Enable Generator

The clock enable generator generates enable signals for clocking the contone CMY/watermark pixel input, the black dot input, and the CMY dot output.

As described earlier, the contone pixel input buffer is used as both a line buffer and a FIFO. Each line is read once and then used *contone scale factor* times. FIFO write address wrapping is disabled until the start of the final replicated use of the line, at which time the clock enable generator generates a *contone line advance enable* signal which enables wrapping.

The clock enable generator also generates an *even* signal which is used to select the even or odd set of output dot FIFOs, and a *margin* signal which is used to generate white dots when the current dot position is in the left or right margin of the page.

The clock enable generator uses a set of counters. The internal logic of the counters is defined in Table 7. The logic of the clock enable signals is defined in Table 8.

Table 7. Clock enable generator counter logic

counter	abbr.	w.	data	load condition	decrement condition
dot	D	14	page width	$RP^a \vee EOL^b$	$(D > 0) \wedge clk$
line	L	15	page length	RP	$(L > 0) \wedge EOL$
left margin	LM	14	left margin	$RP \vee EOL$	$(LM > 0) \wedge clk$
right margin	RM	14	right margin	$RP \vee EOL$	$(RM > 0) \wedge clk$
even/odd dot	E	1	0	$RP \vee EOL$	clk
black dot	BD	14	black width	$RP \vee EOL$	$(LM = 0) \wedge (BD > 0) \wedge clk$
contone dot	CD	14	contone width	$RP \vee EOL$	$(LM = 0) \wedge (CD > 0) \wedge clk$
contone sub-pixel	CSP	4	contone scale factor	$RP \vee EOL \vee (CSP = 0)$	$(LM = 0) \wedge clk$
contone sub-line	CSL	4	contone scale factor	$RP \vee (CSL = 0)$	$EOL \wedge clk$

a. RP (reset page) condition: external signal

b. EOL (end-of-line) condition: $(D=0) \wedge (BD=0) \wedge (CD=0)$

Table 8. Clock enable generator output signal logic

output signal	condition
output dot clock enable	$(D > 0) \wedge \neg EOP^a$
black dot clock enable	$(LM = 0) \wedge (BD > 0) \wedge \neg EOP$
contone pixel clock enable	$(LM = 0) \wedge (CD > 0) \wedge (CSP = 0) \wedge \neg EOP$
contone line advance enable	$(CSL = 0) \wedge \neg EOP$
even	$E = 0$
margin	$(LM = 0) \vee (RM = 0)$

a. EOP (end-of-page) condition: $L = 0$

12.2 PRINTHEAD INTERFACE

Netpage Printer uses an 8½" CMYK Memjet printhead, as described in Section 14. The printhead consists of 17 segments arranged in 2 segment groups. The first segment group contains 9 segments, and the second group contains 8 segments. There are 13,600 nozzles of each color in the printhead, making a total of 54,400 nozzles.

The printhead interface is a standard Memjet printhead interface, as described in Section 15, configured with the following operating parameters:

- MaxColors = 4
- SegmentsPerXfer = 9
- SegmentGroups = 2

Although the printhead interface has a number of external connections, not all used for an 8½" printhead, so not all are connected to external pins on the print engine. Specifically, the value for SegmentGroups implies that there are only 2 SRClock pins and 2 SenseSegSelect pins. All 36 ColorData pins are required, however.

12.3 POSITION TAG ENCODER

The position tag encoder encodes the page id of the page being printed, together with the current x-y position on the page, into an error-correctably encoded position tag which is subsequently printed in infrared (IR) ink on the page.

The position tag encoder takes the following as input:

- the page id
- a pressure sensitive area bitmap at tag resolution
- an active area bitmap at tag resolution

It writes a bi-level IR bitstream to the bi-level IR FIFO. The position tag encoder design as defined here allows a single page id per printed page, and provides for a page size of up to 21.7 inches per dimension ($17 \times 4 \times 512 / 1600$). In addition, the position tag encoder allows for both landscape and portrait orientations.

The position tag encoder consists of two stages connected by a buffer. The first stage, the *tag data encoder*, Reed-Solomon encodes the data to be placed in the tags. The second stage, the *tag formatter*, places the encoded data into the tag format and passes it on to the bi-level IR FIFO. Since the tag formatter runs one line of tags behind the tag data encoder, the tag data buffer is in fact a double-buffer. The relationship is shown in Figure 32.

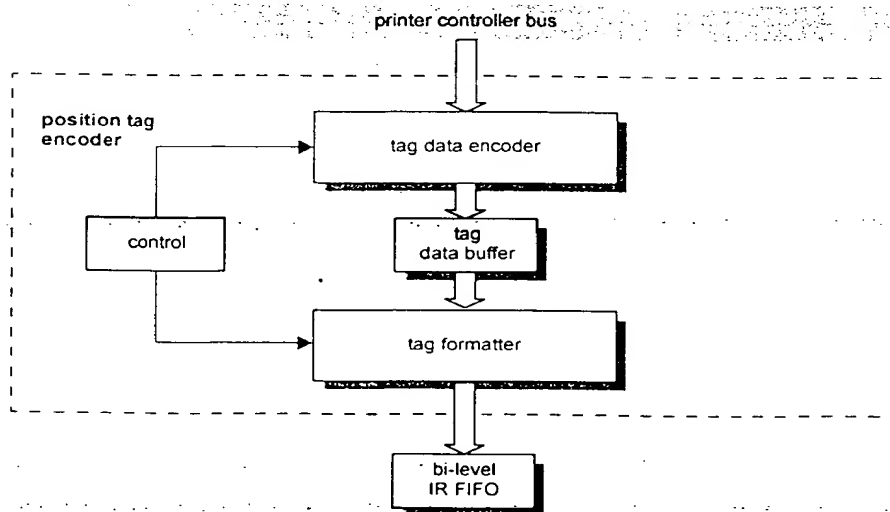


Figure 32. Two stages of position tag encoder

The position tag encoder provides full macrodot-to-dot conversion, thereby eliminating any special knowledge of macrodots from the bi-level FIFO user. However, the position tag encoder itself does make use of the redundancy caused by dot replication in macrodot

formation with respect to the dot axis only. Since the macrodot-to-dot ratio is 1:4, the position tag encoder operates at 1/4 of the dot frequency (60 MHz).

12.3.1 Tag Data Buffer

The tag data buffer consists of two identical buffers. One is read by the tag formatter while the other is being written to by the tag data encoder. The two buffers are then logically swapped.

Each buffer contains $256T$ bits, where T is the maximum number of tags in a given printed line of dots. The maximum amount of memory required for each buffer is 16KBytes (when $T = 512$). To support a 12-inch printhead, $T = 283$, so the memory required for each buffer is 9056 bytes (55% of the maximum). Note that the remainder of the position tag encoder does not change in size based on T .

As the encoded bits are generated by the tag data encoder, they are written to consecutive bits in the tag data buffer. The data bits are read out in random-access fashion by the tag formatter. Having 256 bits allocated to each tag allows the tag number to be used for generating the high bits of the address and the low 8 bits of the address to come from the sub-tag address generator.

The tag data buffer therefore has the structure as shown by Figure 33:

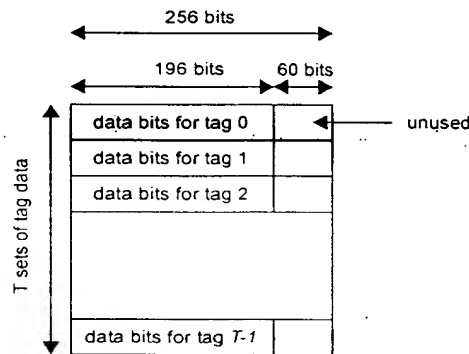


Figure 33. Tag data buffer logical structure

12.3.2 Tag Data Encoder

It is the responsibility of the tag data encoder to produce the 28 bits of encoded mode data and 168 bits of encoded tag data for a total of 196 bits of encoded data. The encoding scheme for position tags is defined in [5].

Each tag is represented on the page by a $17 \times 17 \times 4 \times 4$ structure of dots. However, since the tag data formatter is running at only 1/4 of the dot frequency, this results in a time of 1156 cycles ($17 \times 17 \times 4$) per tag.

The tag data encoder and control block functions are provided by a simple microprocessor core running at 60 MHz. This is desirable compared to the design effort to implement a specific Reed-Solomon encoder for two types of encoding and specific tag codeword construction.

The 196 bits of encoded tag data are written to one of the tag data buffers while the other is being read from by the tag formatter function.

The coordinates encoded within the tags for a given line of tags will depend on the width of the printhead and whether or not the pages are being printed in portrait or landscape mode, as shown in Figure 34.




Figure 34. Relationship of coordinates to tags in a tag line

- If processing a *portrait* page, each printed line corresponds to changes in x coordinates. A printed line of tags therefore has a constant y coordinate, a start x coordinate of 0, and end x coordinate of $T-1$ with a Δx of 1. To advance to the next line of tags, the y coordinate increases by 1 and x is reset to 0.
- If processing a *landscape* page, each printed line corresponds to changes in y coordinates. A printed line of tags therefore has a constant x coordinate, a start y coordinate of 0, and an end y coordinate of $T-1$ with a Δy of 1. To advance to the next line of tags, the x coordinate is decreased by 1 and y is reset to 0.

The high level control function for coordinate generation is outlined here in pseudocode.

```

if portraitMode
  For tagline = 0 to maxTagLines
    For tag = 0 to lastTagInLine
      EncodeTag(tag, pageId, tag, tagLine)
    EndFor
  EndFor
else
  tagX = maxTagLines
  For tagline = 0 to maxTagLines
    For tag = 0 to lastTagInLine
      EncodeTag(tag, pageId, tagX, tag)
      tagX = tagX - 1
    EndFor
  EndFor
EndIf

```

In relation to the EncodeTag function, an examination of the tag structure reveals:

- of the 168 data bits, 90 bits are copied to a maximum of 108 bits in the output with only trivial change. The remaining bits are generated check words based on GF(64).

- of the 28 mode bits, 8 bits are directly copied to the output. The remaining 20 are generated check words based on GF(16).

Therefore the 8 mode bits output and at least 90 bits of the data bits output can be output at a rate approximating 1 cycle per bit. The remaining 70 bits must be generated using Reed-Solomon encoding. This allows approximately 1000 cycles for 70 bits, or approximately 14 cycles per encoded bit. Encoding over GF(16) and GF(64) can be accomplished using small tables totalling 64 and 256 bytes respectively.

The bits are stored in the tag data buffer in bit generation order in order to simplify the writing process. The reading process (in the tag formatter) has random access to the generated bits, but has specific addressing hardware to assist in this task. The order for writing bits is as follows:

- 90-108 bits of data (15-18 \times 6-bit codewords, depending on whether 1s or 0s had to be inserted to generate non 000000 and non 111111 codewords) based on the 90 bits of data (64 bits of pageId, 9 bits of tagX, 9 bits of tagY, 1 bit from the active area bitmap at tag resolution, 1 bit from the pressure sensitive area bitmap at tag resolution and 6 reserved bits).
- 78-60 bits of check words (13-10 \times 6-bit check words, depending on how many data codewords were written)
- 8 bits of mode data (2 \times 4-bit codewords)
- 20 bits of check words (5 \times 4-bit check words)

12.3.3 Tag Formatter

The tag formatter is responsible for merging the encoded tag data with the tag structure, and placing the dots in the IR buffer in the correct order for printing. The encoded tag data is read from the tag data buffer as previously generated by the tag data encoder. The formatted dots are placed in the bi-level IR buffer such that the same data is clocked in 4 times. This allows the tag formatter to run at 1/4 of the dot rate.

We use a simple set of counters for formatting a set of 68 lines of dots (a set of 68 lines equates to 17 rows of 4-dot macrodots). The logic for the 68 lines of dots can be repeated until the page has finished printing. The within-tag counters index into an address table of 9-bit entries with the following attributes:

- If the entry's high bit = 0, then the macrodot comes from an encoded data bit. The encoded data bit to use is given by the concatenation of the current tag# (10 bits) and 8 bits of address from the low 8-bits of the entry.
- If the entry's high bit = 1, then the macrodot comes from the constant structure of a position tag (the bull's-eye and orientation bits). The entry's low bit determines the state of the macrodot.

The address table defines the structure of a positional tag at the macrodot level. It therefore has a constant size of 289×9 -bit entries. The stored 8-bit addresses simply combine the tag encoding structure with the order that the bits are written to the tag data buffer. For the constant part of the tag (the bull's-eye and orientation bits), the entries are simply 100000000 for a white dot and 100000001 for a black dot.

The tag formatter is shown in pseudocode form:

```
For line = 0 to 16
```

```

For dotY = 0 to 3
  For tag = 0 to numTagsInLine
    For pixel = 0 to 16
      adr = TableLookup[line, pixel]
      If (adr5=0)
        bit = TagDataBuffer[tag | adr0-4]
      Else
        bit = adr0
      EndIf
      Place 4 copies of bit in FIFO
    EndFor
  EndFor
EndFor
Swap TagDataBuffers

```

The pseudocode is readily transferred to logic, as illustrated by Figure 35.

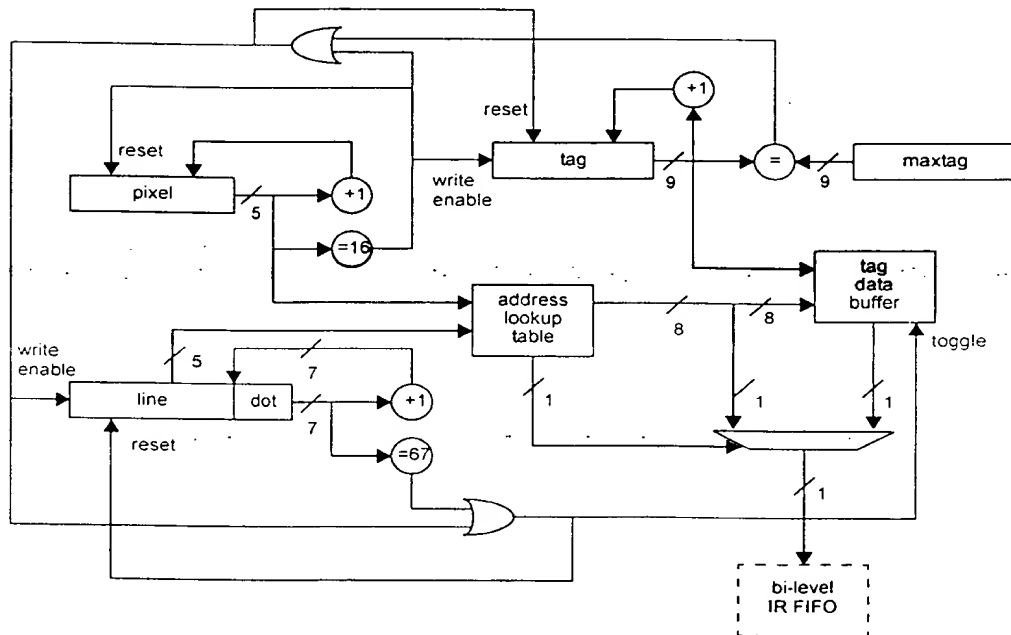


Figure 35. Tag formatter

13 Print Engine Driver

This section discusses rasterizing pages to the internal compressed page format expected by the print engine, as defined in Section 11, in terms of the compressed page driver which hides this device-dependent behavior from the higher-level raster image processor (RIP). The relationship between the RIP, the graphics system, the compressed page driver, and the print engine, is illustrated in Figure 36.




Figure 36. RIP and print engine driver

The RIP and the graphics system are considered generic components and are not defined further at this stage.

13.1 GRAPHICS AND IMAGING MODEL

The compressed page driver is closely coupled with the graphics system used by the RIP, so that the driver can provide device-specific handling for different graphics and imaging operations, in particular compositing operations and text operations.

The graphics system renders images and graphics to a nominal resolution specified by the compressed page driver, but allows the compressed page driver to take control of rendering text. In particular, the graphics system provides sufficient information to the compressed page driver to allow it to render and position text at a higher resolution than the nominal device resolution.

The host graphics system requires random access to a contone page buffer at the nominal device resolution, into which it composites graphics and imaging objects, but it allows the compressed page driver to take control of the actual compositing - i.e. the compressed page driver manages the page buffer.

13.2 TWO-LAYER PAGE BUFFER

The compressed page format contains a 267 ppi contone layer and an 800 dpi black layer. The black layer is conceptually *above* the contone layer, i.e. the black layer is composited *over* the contone layer by the print engine. The compressed page driver therefore main-

tains a page buffer which correspondingly contains a medium-resolution contone layer and a high-resolution black layer.

The graphics systems renders and composites objects into the page buffer bottom-up - i.e. later objects obscure earlier objects. This works naturally when there is only a single layer, but not when there are two layers which will be composited later. It is therefore necessary to detect when an object being placed on the contone layer obscures something on the black layer.

When obscuration is detected, the obscured black pixels are composited with the contone layer and removed from the black layer. The obscuring object is then laid down on the contone layer, possibly interacting with the black pixels in some way. If the compositing mode of the obscuring object is such that no interaction with the background is possible, then the black pixels can simply be discarded without being composited with the contone layer. In practice, of course, there is little interaction between the contone layer and the black layer, since images and text rarely overlap.

The compressed page driver specifies a nominal page resolution of 267 ppi to the graphics system. Where possible the compressed page driver relies on the graphics system to render image and graphics objects to the pixel level at 267 ppi, with the exception of *black* text. The compressed page driver fields all text rendering requests, detects and renders black text at 800 dpi, but returns non-black text rendering requests to the graphics system for rendering at 267 ppi.

Ideally the graphics system and the compressed page driver manipulate color in device-independent RGB, deferring conversion to device-specific CMY until the page is complete and ready to be sent to the printer. This reduces page buffer requirements and makes compositing more rational. Compositing in a device-dependent color space is not ideal.

Ultimately the graphics system asks the compressed page driver to composite each rendered object into the compressed page driver's page buffer. Each such object uses 24-bit contone RGB, and has an explicit (or implicitly opaque) opacity channel.




Figure 37. Two-layer page buffer

The compressed page driver maintains the two-layer page buffer in three parts. The first part is the medium-resolution (267 ppi) contone layer. This consists of a 24-bit RGB bitmap. The second part is a medium-resolution black layer. This consists of an 8-bit opacity bitmap. The third part is a high-resolution (800 dpi) black layer. This consists of a 1-bit opacity bitmap. The medium-resolution black layer is a subsampled version of the high-resolution opacity layer. In practice, assuming the medium resolution is an integer factor n of the high resolution (e.g. $n = 800 / 267 = 3$), each medium-resolution opacity

value is obtained by averaging the corresponding $n \times n$ high-resolution opacity values. This corresponds to box-filtered subsampling. The subsampling of the black pixels effectively antialiases edges in the high-resolution black layer, thereby reducing ringing artifacts when the contone layer is subsequently JPEG-compressed and decompressed.

The structure and size of the page buffer is illustrated in Figure 37.

13.3 COMPOSITING MODEL

For the purposes of discussing the page buffer compositing model, we define the following variables.

Table 9. Compositing variables

variable	description	resolution	format
n	medium to high resolution scale factor	-	-
C_C	contone layer color	medium	8-bit color component
C_G	contone object color	medium	8-bit color component
α_G	contone object opacity	medium	8-bit opacity
α_L	medium-resolution black layer opacity	medium	8-bit opacity
α_B	black layer opacity	high	1-bit opacity
α_T	black object opacity	high	1-bit opacity

When a black object of opacity α_T is composited with the black layer, the black layer is updated as follows:

$$\alpha_{B_{u,v}} = \alpha_{B_{u,v}} \vee \alpha_{T_{u,v}} \quad (1)$$

$$\alpha_{L_{u,v}} = \frac{1}{n^2} \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} 255 \alpha_{B_{u+i, v+j}} \quad (2)$$

The object opacity is simply *ored* with the black layer opacity (Eq. 1), and the corresponding part of the medium-resolution black layer is re-computed from the high-resolution black layer (Eq. 2).

When a contone object of color C_G and opacity α_G is composited with the contone layer, the contone layer and the black layer are updated as follows:

$$C_{C_{u,v}} = C_{C_{u,v}} (1 - \alpha_{L_{u,v}}) \text{ if } (\alpha_{G_{u,v}} > 0) \quad (3)$$

$$\alpha_{L_{u,v}} = 0 \text{ if } (\alpha_{G_{u,v}} > 0) \quad (4)$$

$$\alpha_{B_{u,v}} = 0 \text{ if } (\alpha_{G_{u,v}/n, v/n} > 0) \quad (5)$$

$$C_{C_{u,v}} = C_{C_{u,v}} (1 - \alpha_{G_{u,v}}) + C_{G_{u,v}} \alpha_{G_{u,v}} \quad (6)$$

Wherever the contone object hides the black layer, even if not fully opaquely, the affected black layer pixels are composited with the contone layer (Eq. 3) and removed from the

black layer (Eq. 4 and Eq. 5). The contone object is then composited with the contone layer (Eq. 6).

13.4 PAGE COMPRESSION

Once page rendering is complete, the compressed page driver converts the contone layer to printer-specific CMY with the help of color management functions provided by the graphics system.

The compressed page driver then compresses and packages the black layer and the contone layer into the compressed page format defined in Section 11. This is subsequently expanded in real time by the print engine.

The forward discrete cosine transform (DCT) is the costliest part of JPEG compression. In current high-quality software implementations, the forward DCT of each 8×8 block requires 12 integer multiplications and 32 integer additions [10]. This places a trivial additional load on each RIP DSP.

MEMJET PRINTHEAD

14 Memjet Printhead

A Memjet printhead is a drop-on-demand 1600 dpi inkjet printer that produces bi-level dots in up to 4 colors to produce a printed page of a particular width. Since the printhead prints dots at 1600 dpi, each dot is approximately $22.5\mu\text{m}$ in diameter, and spaced $15.875\mu\text{m}$ apart. Because the printing is bi-level, the input image should be dithered or error-diffused for best results.

Typically a Memjet printhead for a particular application is page-width. This enables the printhead to be stationary and allows the paper to move past the printhead. Figure 38 illustrates a typical configuration.




Figure 38. A Memjet printhead

A Memjet printhead is composed of a number of identical 1/2 inch Memjet segments. The segment is therefore the basic building block for constructing a printhead.

14.1 THE STRUCTURE OF A MEMJET SEGMENT

This section examines the structure of a single segment, the basic building block for constructing Memjet printheads.

14.1.1 Grouping of Nozzles Within a Segment

The nozzles within a single segment are grouped for reasons of physical stability as well as minimization of power consumption during printing. In terms of physical stability, a total of 10 nozzles share the same ink reservoir. In terms of power consumption, groupings are made to enable a low-speed and a high-speed printing mode.

Memjet segments support two printing speeds to allow speed/power consumption trade-offs to be made in different product configurations.

In the low-speed printing mode, 4 nozzles of each color are fired from the segment at a time. The exact number of nozzles fired depends on how many colors are present in the printhead. In a four color (e.g. CMYK) printing environment this equates to 16 nozzles firing simultaneously. In a three color (e.g. CMY) printing environment this equates to 12 nozzles firing simultaneously. To fire all the nozzles in a segment, 200 different sets of nozzles must be fired.

In the high-speed printing mode, 8 nozzles of each color are fired from the segment at a time. The exact number of nozzles fired depends on how many colors are present in the printhead. In a four color (e.g. CMYK) printing environment this equates to 32 nozzles firing simultaneously. In a three color (e.g. CMY) printing environment this equates to 24 nozzles firing simultaneously. To fire all the nozzles in a segment, 100 different sets of nozzles must be fired.

The power consumption in the low-speed mode is half that of the high-speed mode. Note, however, that the energy consumed to print a page is the same in both cases.

14.1.1.1 10 Nozzles Make a Pod

A single pod consists of 10 nozzles sharing a common ink reservoir. 5 nozzles are in one row, and 5 are in another. Each nozzle produces dots $22.5\mu\text{m}$ in diameter spaced on a $15.875\mu\text{m}$ grid to print at 1600 dpi. Figure 39 shows the arrangement of a single pod, with the nozzles numbered according to the order in which they must be fired.




Figure 39. A single pod, numbered by firing order

Although the nozzles are fired in this order, the relationship of nozzles and physical placement of dots on the printed page is different. The nozzles from one row represent the even dots from one line on the page, and the nozzles on the other row represent the odd dots from the adjacent line on the page. Figure 40 shows the same pod with the nozzles numbered according to the order in which they must be loaded.




Figure 40. A single pod, numbered by load order

The nozzles within a pod are therefore logically separated by the width of 1 dot. The exact distance between the nozzles will depend on the properties of the Memjet firing mechanism. The printhead is designed with staggered nozzles designed to match the flow of paper.

14.1.1.2 1 Pod of Each Color Makes a Chromapod

One pod of each color are grouped together into a *chromapod*. The number of pods in a chromapod will depend on the particular application. In a monochrome printing system (such as one that prints only black), there is only a single color and hence a single pod. Photo printing application printheads require 3 colors (cyan, magenta, yellow), so Memjet segments used for these applications will have 3 pods per chromapod (one pod of each color). The expected maximum number of pods in a chromapod is 4, as used in a CMYK (cyan, magenta, yellow, black) printing system (such as a desktop printer). This maximum of 4 colors is not imposed by any physical constraints - it is merely an expected maximum

from the expected applications (of course, as the number of colors increases the cost of the segment increases and the number of these larger segments that can be produced from a single silicon wafer decreases).

A chromapod represents different color components of the same horizontal set of 10 dots on different lines. The exact distance between different color pods depends on the Memjet operating parameters, and may vary from one Memjet design to another. The distance is considered to be a constant number of dot-widths, and must therefore be taken into account when printing: the dots printed by the cyan nozzles will be for different lines than those printed by the magenta, yellow or black nozzles. The printing algorithm must allow for a variable distance up to about 8 dot-widths between colors. Figure 41 illustrates a single chromapod for a CMYK printing application.




Figure 41. A Single Chromapod Contains 1 Pod of each Color

14.1.1.3 5 Chromapods Make a Podgroup

5 chromapods are organized into a single *podgroup*. A podgroup therefore contains 50 nozzles for each color. The arrangement is shown in Figure 42, with chromapods numbered 0-4 and using a CMYK chromapod as the example. Note that the distance between adjacent chromapods is exaggerated for clarity.




Figure 42. A Single Podgroup Contains 5 Chromapods

14.1.1.4 2 Podgroups Make a Phasegroup

2 podgroups are organized into a single *phasegroup*. The phasegroup is so named because groups of nozzles within a phasegroup are fired simultaneously during a given firing phase (this is explained in more detail below). The formation of a phasegroup from 2 podgroups is entirely for the purposes of low-speed and high-speed printing via 2 PodgroupEnable lines.

During low-speed printing, only one of the two PodgroupEnable lines is set in a given firing pulse, so only one podgroup of the two fires nozzles. During high-speed printing, both PodgroupEnable lines are set, so both podgroups fire nozzles. Consequently a low-speed print takes twice as long as a high-speed print, since the high-speed print fires twice as many nozzles at once.

Figure 43 illustrates the composition of a phasegroup. The distance between adjacent podgroups is exaggerated for clarity.



Figure 43. A single phasegroup contains 2 podgroups

14.1.1.5 2 Phasegroups Make a Firegroup

Two phasegroups (PhasegroupA and PhasegroupB) are organized into a single *firegroup*, with 4 firegroups in each segment. Firegroups are so named because they all fire the same nozzles simultaneously. Two enable lines, AEnable and BEnable, allow the firing of PhasegroupA nozzles and PhasegroupB nozzles independently as different firing phases. The arrangement is shown in Figure 44. The distance between adjacent groupings is exaggerated for clarity.

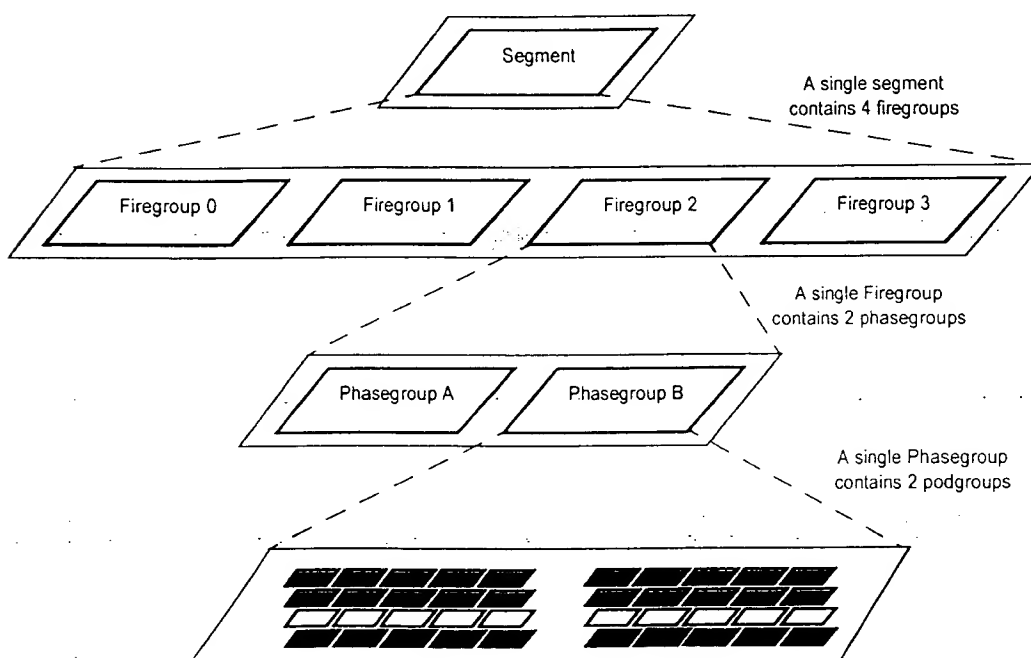


Figure 44. Relationship between Segments, Firegroups, Phasegroups, Podgroups and ChromaPods

14.1.1.6 Nozzle Grouping Summary

Table 10 is a summary of the nozzle groupings in a segment assuming a CMYK chromapod.

Table 10. Nozzle Groupings for a single segment

Name of Grouping	Composition	Replication Ratio	Nozzle Count
Nozzle	Base unit	1:1	1
Pod	Nozzles per pod	10:1	10
Chromapod	Pods per chromapod	C:1	10C
Podgroup	Chromapods per podgroup	5:1	50C
Phasegroup	Podgroups per phasegroup	2:1	100C
Firegroup	Phasegroups per firegroup	2:1	200C
Segment	Firegroups per segment	4:1	800C

The value of C, the number of colors contained in the segment, determines the total number of nozzles.

- With a 4 color segment, such as CMYK, the number of nozzles per segment is 3,200.
- With a 3 color segment, such as CMY, the number of nozzles per segment is 2,400.
- In a monochrome environment, the number of nozzles per segment is 800.

14.1.2 Load and Print Cycles

A single segment contains a total of 800C nozzles, where C is the number of colors in the segment. A *Print Cycle* involves the firing of up to all of these nozzles, dependent on the information to be printed. A *Load Cycle* involves the loading up of the segment with the information to be printed during the subsequent Print Cycle.

Each nozzle has an associated *NozzleEnable* bit that determines whether or not the nozzle will fire during the Print Cycle. The NozzleEnable bits (one per nozzle) are loaded via a set of shift registers.

Logically there are C shift registers per segment (one per color), each 800 deep. As bits are shifted into the shift register for a given color they are directed to the lower and upper nozzles on alternate pulses. Internally, each 800-deep shift register is comprised of two 400-deep shift registers: one for the upper nozzles, and one for the lower nozzles. Alternate bits are shifted into the alternate internal registers. As far as the external interface is concerned however, there is a single 800 deep shift register.

Once all the shift registers have been fully loaded (800 load pulses), all of the bits are transferred in parallel to the appropriate NozzleEnable bits. This equates to a single parallel transfer of 800C bits. Once the transfer has taken place, the Print Cycle can begin. The Print Cycle and the Load Cycle can occur simultaneously as long as the parallel load of all NozzleEnable bits occurs at the end of the Print Cycle.

14.1.2.1 Load Cycle

The Load Cycle is concerned with loading the segment's shift registers with the next Print Cycle's NozzleEnable bits.

Each segment has C inputs directly related to the C shift registers (where C is the number of colors in the segment). These inputs are named *ColorNData*, where N is 1 to C (for example, a 4 color segment would have 4 inputs labeled *Color1Data*, *Color2Data*, *Color3Data* and *Color4Data*). A single pulse on the *SRClock* line transfers C bits into the appropriate shift registers. Alternate pulses transfer bits to the lower and upper nozzles respectively. A total of $800C$ pulses are required for the complete transfer of data. Once all $800C$ bits have been transferred, a single pulse on the *PTransfer* line causes the parallel transfer of data from the shift registers to the appropriate NozzleEnable bits.

The parallel transfer via a pulse on *PTransfer* must take place *after* the Print Cycle has finished. Otherwise the NozzleEnable bits for the line being printed will be incorrect.

It is important to note that the odd and even dot outputs, although printed during the same Print Cycle, do not appear on the same physical output line. The physical separation of odd and even nozzles within the printhead, as well as separation between nozzles of different colors ensures that they will produce dots on different lines of the page. This relative difference must be accounted for when loading the data into the printhead. The actual difference in lines depends on the characteristics of the inkjet mechanism used in the printhead. The differences can be defined by variables D_1 and D_2 where D_1 is the distance between nozzles of different colors, and D_2 is the distance between nozzles of the same color. Table 11 shows the dots transferred to a C color segment on the first 4 pulses.

Table 11. Order of Dots Transferred to a Segment

Pulse	Dot	Color1 Line	Color2 Line	Color3 Line	ColorC Line
1	0	N	$N+D_1^a$	$N+2D_1$	$N+(C-1)D_1$
2	1	$N+D_2^b$	$N+D_1+D_2$	$N+2D_1+D_2$	$N+(C-1)D_1+D_2$
3	2	N	$N+D_1$	$N+2D_1$	$N+(C-1)D_1$
4	3	$N+D_2$	$N+D_1+D_2$	$N+2D_1+D_2$	$N+(C-1)D_1+D_2$

a. D_1 = number of lines between the nozzles of one color and the next (likely = 4-8)

b. D_2 = number of lines between two rows of nozzles of the same color (likely = 1)

And so on for all $800C$ pulses.

Data can be clocked into a segment at a maximum rate of 20 MHz, which will load the entire $800C$ bits of data in $40\mu s$.

14.1.2.2 Print Cycle

A single Memjet printhead segment contains 800 nozzles. To fire them all at once would consume too much power and be problematic in terms of ink refill and nozzle interference. This problem is made more apparent when we consider that a Memjet printhead is composed of multiple 1/2 inch segments, each with 800 nozzles. Consequently two firing modes are defined: a low-speed printing mode and a high-speed printing mode:

- In the low-speed print mode, there are 200 phases, with each phase firing $4C$ nozzles (C per firegroup, where C is the number of colors).
- In the high-speed print mode, there are 100 phases, with each phase firing $8C$ nozzles ($2C$ per firegroup, where C is the number of colors).

The nozzles to be fired in a given firing pulse are determined by

- 3 bits **ChromapodSelect** (select 1 of 5 chromapods from a firegroup)
- 4 bits **NozzleSelect** (select 1 of 10 nozzles from a pod)
- 2 bits of **PodgroupEnable** lines (select 0, 1, or 2 podgroups to fire)

When one of the PodgroupEnable lines is set, only the specified Podgroup's 4 nozzles will fire as determined by ChromapodSelect and NozzleSelect. When both of the PodgroupEnable lines are set, both of the podgroups will fire their nozzles. For the low-speed mode, two fire pulses are required, with PodgroupEnable = 10 and 01 respectively. For the high-speed mode, only one fire pulse is required, with PodgroupEnable = 11.

The duration of the firing pulse is given by the **AEnable** and **BEnable** lines, which fire the PhasegroupA and PhasegroupB nozzles from all firegroups respectively. The typical duration of a firing pulse is 1.3 - 1.8 μ s. The duration of a pulse depends on the viscosity of the ink (dependent on temperature and ink characteristics) and the amount of power available to the printhead. See Section 14.1.3 on page 61 for details on feedback from the printhead in order to compensate for temperature change.

The AEnable and BEnable are separate lines in order that the firing pulses can overlap. Thus the 200 phases of a low-speed Print Cycle consist of 100 A phases and 100 B phases, effectively giving 100 sets of Phase A and Phase B. Likewise, the 100 phases of a high-speed print cycle consist of 50 A phases and 50 B phases, effectively giving 50 phases of phase A and phase B.

Figure 45 shows the AEnable and BEnable lines during a typical Print Cycle. In a high-speed print there are 50 2μ s cycles, while in a low-speed print there are 100 2μ s cycles.




Figure 45. AEnable and BEnable During a Typical Print Cycle

For the high-speed printing mode, the firing order is:

- ChromapodSelect 0, NozzleSelect 0, PodgroupEnable 11 (Phases A and B)
- ChromapodSelect 1, NozzleSelect 0, PodgroupEnable 11 (Phases A and B)
- ChromapodSelect 2, NozzleSelect 0, PodgroupEnable 11 (Phases A and B)
- ChromapodSelect 3, NozzleSelect 0, PodgroupEnable 11 (Phases A and B)
- ChromapodSelect 4, NozzleSelect 0, PodgroupEnable 11 (Phases A and B)
- ChromapodSelect 0, NozzleSelect 1, PodgroupEnable 11 (Phases A and B)
- ...
- ChromapodSelect 3, NozzleSelect 9, PodgroupEnable 11 (Phases A and B)
- ChromapodSelect 4, NozzleSelect 9, PodgroupEnable 11 (Phases A and B)

For the low-speed printing mode, the firing order is similar. For each phase of the high speed mode where PodgroupEnable was 11, two phases of PodgroupEnable = 01 and 10 are substituted as follows:

- ChromapodSelect 0, NozzleSelect 0, PodgroupEnable 01 (Phases A and B)
- ChromapodSelect 0, NozzleSelect 0, PodgroupEnable 10 (Phases A and B)
- ChromapodSelect 1, NozzleSelect 0, PodgroupEnable 01 (Phases A and B)
- ChromapodSelect 1, NozzleSelect 0, PodgroupEnable 10 (Phases A and B)
- ...
- ChromapodSelect 3, NozzleSelect 9, PodgroupEnable 01 (Phases A and B)
- ChromapodSelect 3, NozzleSelect 9, PodgroupEnable 10 (Phases A and B)
- ChromapodSelect 4, NozzleSelect 9, PodgroupEnable 01 (Phases A and B)
- ChromapodSelect 4, NozzleSelect 9, PodgroupEnable 10 (Phases A and B)

When a nozzle fires, it takes approximately $100\mu\text{s}$ to refill. The nozzle cannot be fired before this refill time has elapsed. This limits the fastest printing speed to $100\mu\text{s}$ per line. In the high-speed print mode, the time to print a line is $100\mu\text{s}$, so the time between firing a nozzle from one line to the next matches the refill time. The low-speed print mode is slower than this, so is also acceptable.

The firing of a nozzle also causes acoustic perturbations for a limited time within the common ink reservoir of that nozzle's pod. The perturbations can interfere with the firing of another nozzle within the same pod. Consequently, the firing of nozzles within a pod should be offset from each other as long as possible. We therefore fire four nozzles from a **chromapod (one nozzle per color)** and then move onto the next chromapod within the pod-group.

- In the low-speed printing mode the podgroups are fired separately. Thus the 5 chromapods within both podgroups must all fire before the first chromapod fires again, totalling $10 \times 2\mu\text{s}$ cycles. Consequently each pod is fired once per $20\mu\text{s}$.
- In the high-speed printing mode, the podgroups are fired together. Thus the 5 chromapods within a single podgroups must all fire before the first chromapod fires again, totalling $5 \times 2\mu\text{s}$ cycles. Consequently each pod is fired once per $10\mu\text{s}$.

As the ink channel is $300\mu\text{m}$ long and the velocity of sound in the ink is around 1500m/s , the resonant frequency of the ink channel is 2.5MHz . Thus the low-speed mode allows 50 resonant cycles for the acoustic pulse to dampen, and the high-speed mode allows 25 resonant cycles. Consequently any acoustic interference is minimal in both cases.

14.1.3 Feedback from a Segment

A segment produces several lines of feedback. The feedback lines are used to adjust the timing of the firing pulses. Since multiple segments are collected together into a printhead, it is effective to share the feedback lines as a tri-state bus, with only one of the segments placing the feedback information on the feedback lines.

A pulse on the segment's *SenseSegSelect* line ANDed with data on *Color1Data* selects if the particular segment will provide the feedback. The feedback sense lines will come from that segment until the next *SenseSegSelect* pulse. The feedback sense lines are as follows:

- *Tsense* informs the controller how hot the printhead is. This allows the controller to adjust timing of firing pulses, since temperature affects the viscosity of the ink.
- *Vsense* informs the controller how much voltage is available to the actuator. This allows the controller to compensate for a flat battery or high voltage source by adjusting the pulse width.
- *Rsense* informs the controller of the resistivity (Ohms per square) of the actuator heater. This allows the controller to adjust the pulse widths to maintain a constant energy irrespective of the heater resistivity.
- *Wsense* informs the controller of the width of the critical part of the heater, which may vary up to $\pm 5\%$ due to lithographic and etching variations. This allows the controller to adjust the pulse width appropriately.

14.1.4 Preheat Cycle

The printing process has a strong tendency to stay at the equilibrium temperature. To ensure that the first section of the printed photograph has a consistent dot size, the equilibrium temperature must be met *before* printing any dots. This is accomplished via a preheat cycle.

The Preheat cycle involves a single Load Cycle to all nozzles of a segment with *Is* (i.e. setting all nozzles to fire), and a number of short firing pulses to each nozzle. The duration of the pulse must be insufficient to fire the drops, but enough to heat up the ink. Altogether about 200 pulses for each nozzle are required, cycling through in the same sequence as a standard Print Cycle.

Feedback during the Preheat mode is provided by *Tsense*, and continues until equilibrium temperature is reached (about 30° C above ambient). The duration of the Preheat mode is around 50 milliseconds, and depends on the ink composition.

Preheat is performed before each print job. This does not affect performance as it is done while the data is being transferred to the printer.

14.1.5 Cleaning Cycle

In order to reduce the chances of nozzles becoming clogged, a cleaning cycle can be undertaken before each print job. Each nozzle is fired a number of times into an absorbent sponge.

The cleaning cycle involves a single Load Cycle to all nozzles of a segment with *Is* (i.e. setting all nozzles to fire), and a number of firing pulses to each nozzle. The nozzles are cleaned via the same nozzle firing sequence as a standard Print Cycle. The number of times that each nozzle is fired depends upon the ink composition and the time that the

printer has ben idle. As with preheat, the cleaning cycle has no effect on printer performance.

14.1.6 Printhead Interface Summary

Each segment has the following connections to the bond pads:

Table 12. Segment Interface Connections

Name	Lines	Description
Chromapod Select	3	Select which chromapod will fire (0-4)
NozzleSelect	4	Select which nozzle from the pod will fire (0-9)
PodgroupEnable	2	Enable the podgroups to fire (choice of: 01, 10, 11)
AEnable	1	Firing pulse for podgroup A
BEnable	1	Firing pulse for podgroup B
ColorNData	C	Input to shift registers (1 bit for each of C colors in the segment)
SRClock	1	A pulse on SRClock (ShiftRegisterClock) loads C bits from Color-Data into the C shift registers.
PTransfer	1	Parallel transfer of data from the shift registers to the internal NozzleEnable bits (one per nozzle).
SenseSegSelect	1	A pulse on SenseSegSelect ANDed with data on Color1Data selects the sense lines for this segment.
Tsense	1	Temperature sense
Vsense	1	Voltage sense
Rsense	1	Resistivity sense
Wsense	1	Width sense
Logic GND	1	Logic ground
Logic PWR	1	Logic power
V-	21	Actuator Ground
V+	21	Actuator Power
TOTAL	62+C	(if C is 4, Total = 66)

14.2 MAKING MEMJET PRINTHEADS OUT OF SEGMENTS

A Memjet printhead is composed of a number of identical 1/2 inch printhead segments. These 1/2 inch segments are manufactured together or placed together after manufacture to produce a printhead of the desired length. Each 1/2 inch segments prints 800 1600 dpi bi-level dots in up to 4 colors over a different part of the page to produce the final image. Although each segment produces 800 dots of the final image, each dot is represented by a combination of colored inks.

A 4-inch printhead, for example, consists of 8 segments, typically manufactured as a monolithic printhead. In a typical 4-color printing application (cyan, magenta, yellow, black), each of the segments prints bi-level cyan, magenta, yellow and black dots over a different part of the page to produce the final image. The positions of the segments are shown in Figure 46.

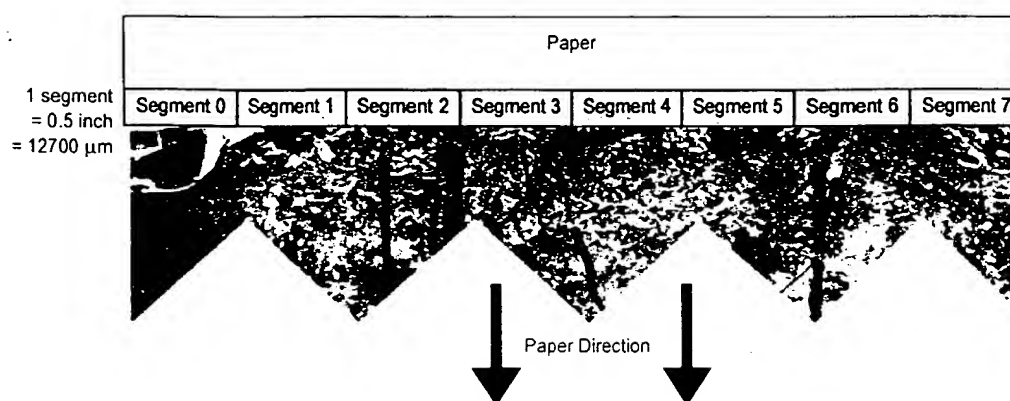


Figure 46. Arrangement of Segments in a 4-inch Printhead

An 8-inch printhead can be constructed from two 4-inch printheads or from a single 8-inch printhead consisting of 16 segments. Regardless of the construction mechanism, the effective printhead is still 8 inches in length.

A 2-inch printhead has a similar arrangement, but only uses 4 segments. Likewise, a full-bleed A4/Letter printer uses 17 segments for an effective 8½" printing area.

Since the total number of nozzles in a segment is 800C (see Table 10), the total number of nozzles in a given printhead with S segments is 800CS. Thus segment N is responsible for printing dots 800N to 800N+799.

A number of considerations must be made when wiring up a printhead. As the width of the printhead increases, the number of segments increases, and the number of connections also increases. Each segment has its own ColorData connections (C of them), as well as SRClock and other connections for loading and printing.

14.2.1 Loading Considerations

When the number of segments S is small it is reasonable to load all the segments simultaneously by using a common SRClock line and placing C bits of data on each of the ColorData inputs for the segments. In a 4-inch printer, $S=8$, and therefore the total number of bits to transfer to the printhead in a single SRClock pulse is 32. However for an 8-inch printer, $S=16$, and it is unlikely to be reasonable to have 64 data lines running from the print data generator to the printhead.

Instead, it is convenient to group a number of segments together for loading purposes. Each group of segments is small enough to be loaded simultaneously, and share an SRClock. For example, an 8-inch printhead can have 2 segment groups, each segment group containing 8 segments. 32 ColorData lines can be shared for both groups, with 2 SRClock lines, one per segment group.

When the number of segment groups is not easily divisible, it is still convenient to group the segments. One example is a 8½" printer for producing A4/Letter pages. There are 17 segments, and these can be grouped as two groups of 9 (9C bits of data going to each segment, with all 9C bits used in the first group, and only 8C bits used for the second group), or as 3 groups of 6 (again, C bits are unused in the last group).

As the number of segment groups increases, the time taken to load the printhead increases. When there is only one group, 800 load pulses are required (each pulse transfers C data bits). When there are G groups, 800G load pulses are required. The bandwidth of the connection between the data generator and the printhead must be able to cope and be within the allowable timing parameters for the particular application.

If G is the number of segment groups, and L is the largest number of segments in a group, the printhead requires LC ColorData lines and G SRClock lines. Regardless of G , only a single PTransfer line is required - it can be shared across all segments.

Since L segments in each segment group are loaded with a single SRClock pulse, any printing process must produce the data in the correct sequence for the printhead. As an example, when $G=2$ and $L=4$, the first SRClock1 pulse will transfer the ColorData bits for the next Print Cycle's dot 0, 800, 1600, and 2400. The first SRClock2 pulse will transfer the ColorData bits for the next Print Cycle's dot 3200, 4000, 4800, and 5600. The second SRClock1 pulse will transfer the ColorData bits for the next Print Cycle's dot 1, 801, 1601, and 2401. The second SRClock2 pulse will transfer the ColorData bits for the next Print Cycle's dot 3201, 4001, 4801 and 5601.

After 800G SRClock pulses (800 to each of SRClock1 and SRClock2), the entire line has been loaded into the printhead, and the common PTransfer pulse can be given.

It is important to note that the odd and even color outputs, although printed during the same Print Cycle, do not appear on the same physical output line. The physical separation of odd and even nozzles within the printhead, as well as separation between nozzles of different colors ensures that they will produce dots on different lines of the page. This relative difference must be accounted for when loading the data into the printhead. The actual difference in lines depends on the characteristics of the inkjet mechanism used in the printhead. The differences can be defined by variables D_1 and D_2 where D_1 is the distance between nozzles of different colors, and D_2 is the distance between nozzles of the same

color. Considering only a single segment group, Table 13 shows the dots transferred to segment n of a printhead during the first 4 pulses of the shared SRClock.

Table 13. Order of Dots Transferred to a Segment in a Printhead

Pulse	Dot	Color1 Line	Color2 Line	ColorC Line
1	$800S^a$	N	$N+D_1^b$	$N+(C-1)D_1$
2	$800S+1$	$N+D_2^c$	$N+D_1+D_2$	$N+(C-1)D_1+D_2$
3	$800S+2$	N	$N+D_1$	$N+(C-1)D_1$
4	$800S+3$	$N+D_2$	$N+D_1+D_2$	$N+(C-1)D_1+D_2$

a. S = segment number

b. D_1 = number of lines between the nozzles of one color and the next (likely = 4-8)

c. D_2 = number of lines between two rows of nozzles of the same color (likely = 1)

And so on for all 800 SRClock pulses to the particular segment group.

14.2.2 Printing Considerations

With regards to printing, we print 4C nozzles from each segment in the low-speed printing mode, and 8C nozzles from each segment in the high speed printing mode.

While it is certainly possible to wire up segments in any way, this document only considers the situation where all segments fire simultaneously. This is because the low-speed printing mode allows low-power printing for small printheads (e.g. 2-inch and 4-inch), **and the controller chip design assumes there is sufficient power available for the large print sizes (such as 8-18 inches).** It is a simple matter to alter the connections in the printhead to allow grouping of firing should a particular application require it.

When all segments are fired at the same time 4CS nozzles are fired in the low-speed printing mode and 8CS nozzles are fired in the high-speed printing mode. Since all segments print simultaneously, the printing logic is the same as defined in Section 14.1.2.2 on page 58.

The timing for the two printing modes is therefore:

- 200 μ s to print a line at low speed (comprised of 100 2 μ s cycles)
- 100 μ s to print a line at high speed (comprised of 50 2 μ s cycles)

14.2.3 Feedback Considerations

A segment produces several lines of feedback, as defined in Section 14.1.3 on page 61. The feedback lines are used to adjust the timing of the firing pulses. Since multiple segments are collected together into a printhead, it is effective to share the feedback lines as a tri-state bus, with only one of the segments placing the feedback information on the feedback lines at a time.

Since the selection of which segment will place the feedback information on the shared Tsense, Vsense, Rsense, and Wsense lines uses the ColorI Data line, the groupings of segments for loading data can be used for selecting the segment for feedback.

Just as there are G SRClock lines (a single line is shared between segments of the same segment group), there are G SenseSegSelect lines shared in the same way. When the correct SenseSegSelect line is pulsed, the segment of that group whose Color1Data bit is set will start to place data on the shared feedback lines. The segment previously active in terms of feedback must also be disabled by having a 0 on its Color1Data bit, and this segment may be in a different segment group. Therefore when there is more than one segment group, changing the feedback segment requires two steps: disabling the old segment, and enabling the new segment.

14.2.4 Printhead Connection Summary

This section assumes that a printhead has been constructed from a number of segments as described in the previous sections. It assumes that for data loading purposes, the segments have been grouped into G segment groups, with L segments in the largest segment group. It assumes there are C colors in the printhead. It assumes that the firing mechanism for the printhead is that all segments fire simultaneously, and only one segment at a time places feedback information on a common tri-state bus. Assuming all these things, Table 14 lists the external connections that are available from a printhead:

Table 14. Printhead Connections

Name	#Pins	Description
ChromapodSelect	3	Select which chromapod will fire (0-4)
NozzleSelect	4	Select which nozzle from the pod will fire (0-9)
PodgroupEnable	2	Enable the podgroups to fire (choice of: 01, 10, 11)
AEnable	1	Firing pulse for phasegroup A
BEnable	1	Firing pulse for phasegroup B
ColorData	CL	Inputs to C shift registers of segments 0 to $L-1$
SRClock	G	A pulse on SRClock[N] (ShiftRegisterClock N) loads the current values from ColorData lines into the L segments in segment group N.
PTransfer	1	Parallel transfer of data from the shift registers to the internal NozzleEnable bits (one per nozzle).
SenseSegSelect	G	A pulse on SenseSegSelect N ANDed with data on Color1Data[n] selects the sense lines for segment n in segment group N.
Tsense	1	Temperature sense
Vsense	1	Voltage sense
Rsense	1	Resistivity sense
Wsense	1	Width sense
Logic GND	1	Logic ground
Logic PWR	1	Logic power
V-	Bus bars	Actuator Ground
V+		Actuator Power
TOTAL	18+2G+CL	

15 Memjet Printhead Interface

The printhead interface (PHI) is the means by which the processor loads the Memjet printhead with the dots to be printed, and controls the actual dot printing process. The PHI contains:

- a LineSyncGen unit (LSGU), which provides synchronization signals for multiple chips (allows side-by-side printing and front/back printing) as well as stepper motors.
- a Memjet interface (MJI), which transfers data to the Memjet printhead, and controls the nozzle firing sequences during a print.
- a line loader/format unit (LLFU) which loads the dots for a given print line into local buffer storage and formats them into the order required for the Memjet printhead.

The units within the PHI are controlled by a number of registers that are programmed by the processor. In addition, the processor is responsible for setting up the appropriate parameters in the DMA controller for the transfers from memory to the LLFU. This includes loading white (all 0's) into appropriate colors during the start and end of a page so that the page has clean edges.

The PHI is capable of dealing with a variety of printhead lengths and formats. In terms of broad operating customizations, the PHI is parameterized as follows:

Table 15. Basic Printing Parameters

Name	Description	Range
MaxColors	No of Colors in printhead	1-4
SegmentsPerXfer	No of segments written to per transfer. Is equal to the number of segments in the largest segment group	1-9
SegmentGroups	No of segment groups in printhead	1-4

The internal structure of the PHI allows for a maximum of 4 colors, 9 segments per transfer, and 4 transfers. Transferring 4 colors to 9 segments is 36 bits per transfer, and 4 transfers to 9 segments equates to a maximum printed line length of 18 inches. The total number of dots per line printed by an 18-inch 4 color printhead is 115,200 ($18 \times 1600 \times 4$).

Other example settings are shown in Table 16:

Table 16. Example Settings for Basic Printing Parameters

Printer Length	Printer Type	MaxColors	SegmentsPerXfer	SegmentGroups	Bits per Xfer	Comments
4 inch CMY	Photo	3	8	1	24	
8 inch CMYK	A4/Letter	4	8	2	32	
8½ inch CMYK	A4/Letter full bleed	4	9	2	36	Last xfer not fully used
12 inch CMYK	A4 long / A3 short	4	8	3	32	
16 inch CMYK		4	8	4	32	
17 inch CMYK	A3 long full bleed	4	9	4	36	Last xfer not fully used
18 inch CMYK		4	9	4	36	

15.1 BLOCK DIAGRAM OF PRINTHEAD INTERFACE

The internal structure of the Printhead Interface is shown in Figure 47.

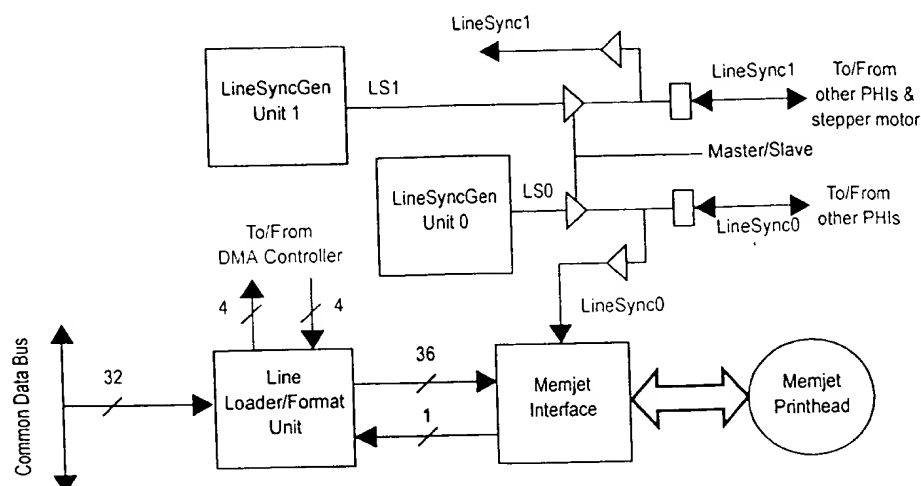


Figure 47. Internal Structure of Printhead Interface

In the PHI there are two LSGUs. The first LSGU produces LineSync0, which is used to control the Memjet Interface in all synchronized chips. The second LSGU produces LineSync1 which is used to pulse the paper drive stepper motor.

The Master/Slave pin on the chip allows multiple chips to be connected together for side-by-side printing, front/back printing etc. via a Master/Slave relationship. When the Master/Slave pin is attached to V_{DD} , the chip is considered to be the Master, and LineSync pulses generated by the two LineSyncGen units are enabled onto the two tri-state LineSync common lines (LineSync0 and LineSync1, shared by all the chips). When the Master/Slave pin is attached to GND, the chip is considered to be the Slave, and LineSync pulses generated by the two LineSyncGen units are not enabled onto the common LineSync lines. In this way, the Master chip's LineSync pulses are used by all PHIs on all the connected chips.

The following sections detail the LineSyncGen Unit, the Line Loader/Format Unit and Memjet Interface respectively.

15.2 LINESYNGEN UNIT

The LineSyncGen units (LSGU) are responsible for generating the synchronization pulses required for printing a page. Each LSGU produces an external LineSync signal to enable line synchronization. The generator inside the LGSU generates a LineSync pulse when told to 'go', and then every so many cycles until told to stop. The LineSync pulse defines the start of the next line.

The exact number of cycles between LineSync pulses is determined by the CyclesBetweenPulses register, one per generator. It must be at least long enough to allow one line to print (100 μ s or 200 μ s depending on whether the speed is low or high) and another line to load, but can be longer as desired (for example, to accommodate special requirements of paper transport circuitry). If the CyclesBetweenPulses register is set to a number less than a line

print time, the page will not print properly since each LineSync pulse will arrive before the particular line has finished printing.

The following interface registers are contained in the LSGU:

Table 17. LineSyncGen Unit Registers

Register Name	Description
CyclesBetweenPulses	The number of cycles to wait between generating one LineSync pulse and the next.
Go	Controls whether the LSGU is currently generating LineSync pulses or not. A write of 1 to this register generates a LineSync pulse, transfers CyclesBetweenPulses to CyclesRemaining, and starts the countdown. When CyclesRemaining hits 0, another LineSync pulse is generated, CyclesBetweenPulses is transferred to CyclesRemaining and the countdown is started again. A write of 0 to this register stops the countdown and no more LineSync pulses are generated.
CyclesRemaining	A status register containing the number of cycles remaining until the next LineSync pulse is generated.

The LineSync pulse is not used directly from the LSGU. The LineSync pulse is enabled onto a tri-state LineSync line only if the Master/Slave pin is set to Master. Consequently the LineSync pulse is only used in the form as generated by the Master chip (pulses generated by Slave chips are ignored).

15.3 MEMJET INTERFACE

The Memjet interface (MJI) transfers data to the Memjet printhead, and controls the nozzle firing sequences during a print.

The MJI is simply a State Machine (see Figure 48) which follows the printhead loading and firing order described in Section 14.2.1 on page 64, Section 14.2.2 on page 65, and includes the functionality of the Preheat Cycle and Cleaning Cycle as described in Section 14.1.4 on page 61 and Section 14.1.5 on page 61. Both high-speed and low-speed printing modes are available, although the MJI always fires a given nozzle from all segments in a printhead simultaneously (there is no separate firing of nozzles from one segment and then others). Dot counts for each color are also kept by the MJI.

The MJI loads data into the printhead from a choice of 2 data sources:

- All 1s. This means that all nozzles will fire during a subsequent Print cycle, and is the standard mechanism for loading the printhead for a preheat or cleaning cycle.
- From the 36-bit input held in the Transfer register of the LLFU. This is the standard means of printing an image. The 36-bit value from the LLFU is directly sent to the printhead and a 1-bit 'Advance' control pulse is sent to the LLFU.

The MJI knows how many lines it has to print for the page. When the MJI is told to 'go', it waits for a LineSync pulse before it starts the first line. Once it has finished loading/printing a line, it waits until the next LineSync pulse before starting the next line. The MJI stops once the specified number of lines has been loaded/printed, and ignores any further LineSync pulses.

The MJJ is therefore directly connected to the LLFU, LineSync0 (shared between all synchronized chips), and the external Memjet printhead. The basic structure is shown in Figure 48.




Figure 48. Memjet Interface

The MJJ accepts 36 bits of data from the LLFU. Of these 36 bits, only the bits corresponding to the number of segments and number of colors will be valid. For example, if there are only 2 colors and 9 segments, bits 0-1 will be valid for segment 0, bits 2-3 will be invalid, bits 4-5 will be valid for segment 1, bits 6-7 will be invalid etc. The state machine does not care which bits are valid and which bits are not valid - it merely passes the bits out to the printhead. The data lines and control signals coming out of the MJJ can be wired appropriately to the pinouts of the chip, using as few pins as required by the application range of the chip (see Section 15.3.1 on page 70 for more information).

15.3.1 Connections to Printhead

The MJJ has a number of connections to the printhead, including a maximum of 4 colors, clocked in to a maximum of 9 segments per transfer to a maximum of 4 segment groups. The lines coming from the MJJ can be directly connected to pins on the chip, although not all lines will always be pins. For example, if the chip is specifically designed for only connecting to 8 inch CMYK printers, only 32 bits of data need to be transferred each transfer pulse. Consequently 32 pins of data out (8 pins per color), and not 36 pins are required. In the same way, only 2 SRClock pulses are required, so only 2 pins instead of 4 pins are required to cater for the different SRClocks. And so on.

If the chip must be completely generic, then all connections from the MJJ must be connected to pins on the chip (and thence to the Memjet printhead).

Table 18 lists the maximum connections from the MJJ, many of which are always connected to pins on the chip. Where the number of pins is variable, a footnote explains what the number of pins depends upon. The sense of input and output is with respect to the MJJ. The names correspond to the pin connections on the printhead.

Table 18. Memjet Interface Connections

Name	#Pins	I/O	Description
Chromapod Select	3	O	Select which chromapod will fire (0-4)
NozzleSelect	4	O	Select which nozzle from the pod will fire (0-9)
PodgroupEnable	2	O	Enable the podgroups to fire (choice of: 01, 10, 11)
AEnable	1	O	Firing pulse for podgroup A. In the current design all segments fire simultaneously, although multiple AEnable lines could be added for dividing the firing sequence over multiple segment groups for reasons of power and speed.

Table 18. Memjet Interface Connections

Name	#Pins	I/O	Description
BEnable	1	O	Firing pulse for podgroup B. In the current design all segments fire simultaneously, although multiple BEnable lines could be added for dividing the firing sequence over multiple segment groups for reasons of power and speed.
Color1Data[0-8]	g ^a	O	Output to Color1Data shift register of segments 0-8
Color2Data[0-8]	g ^b	O	Output to Color2Data shift register of segments 0-8
Color3Data[0-8]	g ^c	O	Output to Color3Data shift register of segments 0-8
Color4Data[0-8]	g ^d	O	Output to Color4Data shift register of segments 0-8
SRClock[1-4]	4 ^e	O	A pulse on SRClock[N] (ShiftRegisterClock) loads the current values from Color1Data[0-8], Color2Data[0-8], Color3Data[0-8] and Color4Data[0-8] into the segment group N on the printhead.
PTransfer	1	O	Parallel transfer of data from the shift registers to the printhead's internal NozzleEnable bits (one per nozzle).
SenseSegSelect[1-4]	4 ^f	O	A pulse on SenseSegSelect[N] ANDed with data on Color1Data[n] enables the sense lines for segment n in segment group N of the printhead.
Tsense	1	I	Temperature sense
Vsense	1	I	Voltage sense
Rsense	1	I	Resistivity sense
Wsense	1	I	Width sense
TOTAL	52		

- Although 9 lines are available from the MJ1, the number of pins coming from the chip will only reflect the actual number of segments in a segment group. The pins for Color1Data are mandatory, since each printhead must print in at least 1 color.
- These lines are only translated into pins if the chip is to control a printhead with at least 2 colors. Although 9 lines are available from the MJ1, the number of pins coming from the chip for Color2Data will only reflect the *actual* number of segments in a segment group.
- These lines are only translated into pins if the chip is to control a printhead with at least 3 colors. Although 9 lines are available from the MJ1, the number of pins coming from the chip for Color3Data will only reflect the *actual* number of segments in a segment group.
- These lines are only translated into pins if the chip is to control a printhead with 4 colors. Although 9 lines are available from the MJ1, the number of pins coming from the chip for Color4Data will only reflect the *actual* number of segments in a segment group.
- Although 4 lines are available from the MJ1, the number of pins coming from the chip will only reflect the *actual* number of segment groups. A minimum of 1 pin is required since there is at least 1 segment group (the entire printhead).
- Although 4 lines are available from the MJ1, the number of pins coming from the chip will only reflect the *actual* number of segment groups. A minimum of 1 pin is required since there is at least 1 segment group (the entire printhead).

15.3.2 Firing Pulse Duration

The duration of firing pulses on the AEnable and BEnable lines depend on the viscosity of the ink (which is dependant on temperature and ink characteristics) and the amount of power available to the printhead. The typical pulse duration range is 1.3 to 1.8 μ s. The MJ1 therefore contains a programmable pulse duration table, indexed by feedback from the printhead. The table of pulse durations allows the use of a lower cost power supply, and aids in maintaining more accurate drop ejection.

The Pulse Duration table has 256 entries, and is indexed by the current Vsense and Tsense settings. The upper 4-bits of address come from Vsense, and the lower 4-bits of address

come from Tsense. Each entry is 8 bits, and represents a fixed point value in the range of 0-4 μ s. The process of generating the AEnable and BEnable lines is shown in Figure 49.

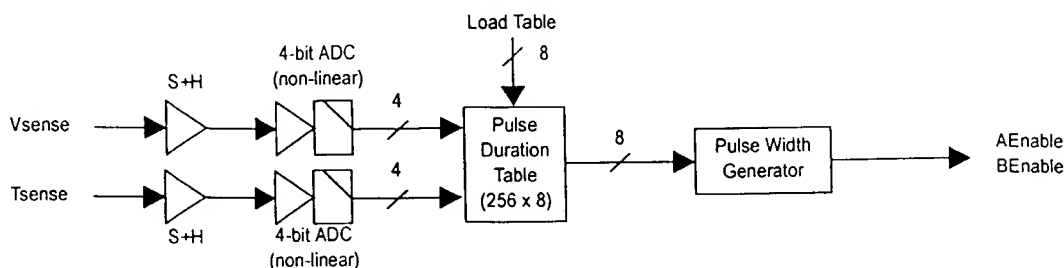


Figure 49. Generation of AEnable and BEnable Pulse Widths

The 256-byte table is written by the CPU before printing the first page. The table may be updated in between pages if desired. Each 8-bit pulse duration entry in the table combines:

- User brightness settings (from the page description)
- Viscosity curve of ink (from the QA Chip)
- Rsense
- Wsense
- Vsense
- Tsense

15.3.3 Dot Counts

The MJ1 maintains a count of the number of dots of each color fired from the printhead. The dot count for each color is a 32-bit value, individually cleared under processor control. At 32-bits length, each dot count can hold a maximum coverage dot count of 17 8-inch \times 12-inch pages, although in typical usage, the dot count will be read and cleared after each page or half-page.

The dot counts are used by the processor to update the QA chip in order to predict when the ink cartridge runs out of ink. The processor knows the volume of ink in the cartridge for each of the colors from the QA chip. Counting the number of drops eliminates the need for ink sensors, and prevents the ink channels from running dry. An updated drop count is written to the QA chip after each page. A new page will not be printed unless there is enough ink left, and allows the user to change the ink without getting a dud half-printed page which must be reprinted.

The layout of the dot counter for Color1 is shown in Figure 50. The remaining 3 dot counters (Color1DotCount, Color2DotCount, and Color3DotCount) are identical in structure.




Figure 50. Dot Count Logic

15.3.4 Registers

The processor communicates with the MJJ via a register set. The registers allow the processor to parameterize a print as well as receive feedback about print progress.

The following registers are contained in the MJJ:

Table 19. Memjet Interface Registers

Register Name	Description
Print Parameters	
SegmentsPerXfer	The number of segments to write to each transfer. This also equals the number of cycles to wait between each transfer (before generating the next Advance pulse). Each transfer has MaxColors × SegmentsPerXfer valid bits.
SegmentGroups	The number of segment groups in the printhead. This equals the number of times that SegmentsPerXfer cycles must elapse before a single dot has been written to each segment of the printhead. The MJJ does this 800 times to completely transfer all the data for the line to the printhead.
PrintSpeed	Whether to print at low or high speed (determines the value on the PodgroupEnable lines during the print).
NumLines	The number of Load/Print cycles to perform.
Monitoring the Print (read only from point of view of processor)	
Status	The Memjet Interface's Status Register
LinesRemaining	The number of lines remaining to be printed. Only valid while Go=1. Starting value is NumLines and counts down to 0.
TransfersRemaining	The number of sets of SegmentGroups transfers remaining before the Printhead is considered loaded for the current line. Starts at 800 and counts down to 0. Only valid while Go=1.
SegGroupsRemaining	The number of segment groups remaining in the current set of transfers of 1 dot to each segment. Starts at SegmentGroups and counts down to 0. Only valid while Go=1.
SenseSegment	The 9-bit value to place on the Color1Data lines during a subsequent feedback SenseSegSelect pulse. Only 1 of the 9 bits should be set, corresponding to one of the (maximum) 9 segments. See SenseSelect for how to determine which of the segment groups to sense.

Table 19. Memjet Interface Registers

Register Name	Description
SetAllNozzles	<p>If non-zero, the 36-bit value written to the printhead during the LoadDots process is all 1s, so that all nozzles will be fired during the subsequent PrintDots process. This is used during the preheat and cleaning cycles.</p> <p>If 0, the 36-bit value written to the printhead comes from the LLFU. This is the case during the actual printing of regular images.</p>
Actions	
Reset	A write to this register resets the MJL, stops any loading or printing processes, and loads all registers with 0.
SenseSelect	<p>A write to this register with any value clears the FeedbackValid bit of the Status register, and the remaining action depends on the values in the LoadingDots and PrintingDots status bits.</p> <p>If either of the status bits are set, the Feedback bit is cleared and nothing more is done.</p> <p>If both status bits are clear, a pulse is given simultaneously on all 4 SenseSegSelect lines with all ColorData bits 0. This stops any existing feedback. Depending on the two low-order bits written to SenseSelect register, a pulse is given on SenseSegSelect1, SenseSegSelect2, SenseSegSelect3, or SenseSegSelect4 line, with the Color1Data bits set according to the SenseSegment register.</p> <p>Once the various sense lines have been tested, the values are placed in the Tsense, Vsense, Rsense, and Wsense registers, and the Feedback bit of the Status register is set.</p>
Go	<p>A write of 1 to this bit starts the LoadDots / PrintDots cycles, which commences with a wait for the first LineSync pulse. A total of NumLines lines are printed, each line being loaded/printed after the receipt of a LineSync pulse. The loading of each line consists of SegmentGroups 36-bit transfers. As each line is printed, LinesRemaining decrements, and TransfersRemaining is reloaded with SegmentGroups again. The status register contains print status information. Upon completion of NumLines, the loading/printing process stops, the Go bit is cleared, and any further LineSync pulses are ignored. During the final print cycle, nothing is loaded into the printhead.</p> <p>A write of 0 to this bit stops the print process, but does not clear any other registers.</p>
ClearCounts	A write to this register clears the Color1DotCount, Color2DotCount, Color3DotCount, and Color4DotCount registers if bits 0, 1, 2, or 3 respectively are set. Consequently a write of 0 has no effect.
Feedback	
Tsense	Read only feedback of Tsense from the last SenseSegSelect pulse sent to segment SenseSegment. Is only valid if the FeedbackValid bit of the Status register is set.
Vsense	Read only feedback of Vsense from the last SenseSegSelect pulse sent to segment SenseSegment. Is only valid if the FeedbackValid bit of the Status register is set.
Rsense	Read only feedback of Rsense from the last SenseSegSelect pulse sent to segment SenseSegment. Is only valid if the FeedbackValid bit of the Status register is set.
Wsense	Read only feedback of Wsense from the last SenseSegSelect pulse sent to segment SenseSegment. Is only valid if the FeedbackValid bit of the Status register is set.
Color1DotCount	Read only 32-bit count of color1 dots sent to the printhead.
Color2DotCount	Read only 32-bit count of color2 dots sent to the printhead.

Table 19. Memjet Interface Registers

Register Name	Description
Color3DotCount	Read only 32-bit count of color3 dots sent to the printhead
Color4DotCount	Read only 32-bit count of color4 dots sent to the printhead

The MJl's Status Register is a 16-bit register with bit interpretations as follows:

Table 20. MJl Status Register

Name	Bits	Description
LoadingDots	1	If set, the MJl is currently loading dots, with the number of dots remaining to be transferred in TransfersRemaining. If clear, the MJl is not currently loading dots
PrintingDots	1	If set, the MJl is currently printing dots. If clear, the MJl is not currently printing dots.
PrintingA	1	This bit is set while there is a pulse on the AEnable line
PrintingB	1	This bit is set while there is a pulse on the BEnable line
FeedbackValid	1	This bit is set while the feedback values Tsense, Vsense, Rsense, and Wsense are valid.
Reserved	3	-
PrintingChromapod	4	This holds the current chromapod being fired while the PrintingDots status bit is set.
PrintingNozzles	4	This holds the current nozzle being fired while the PrintingDots status bit is set.

The following pseudocode illustrates the logic required to load a printhead for a single line. Note that loading commences only after the LineSync pulse arrives. This is to ensure the data for the line has been prepared by the LLFU and is valid for the first transfer to the printhead.

```

Wait for LineSync
For TransfersRemaining = 800 to 0
  For I = 0 to SegmentGroups
    If (SetAllNozzles)
      Set all ColorData lines to be 1
    Else
      Place 36 bit input on 36 ColorData lines
    EndIf
    Pulse SRClock[I]
    Wait SegmentsPerXfer cycles
    Send ADVANCE signal
  EndFor
EndFor

```

15.3.5 Preheat and Cleaning Cycles

The Cleaning and Preheat cycles are simply accomplished by setting appropriate registers in the MJ1:

- SetAllNozzles = 1
- Set the PulseDuration register to either a low duration (in the case of the preheat mode) or to an appropriate drop ejection duration for cleaning mode.
- Set NumLines to be the number of times the nozzles should be fired
- Set the Go bit and then wait for the Go bit to be cleared when the print cycles have completed.

The LSGU must also be programmed to send LineSync pulses at the correct frequency.

15.4 LINE LOADER/FORMAT UNIT

The line loader/format unit (LLFU) loads the dots for a given print line into local buffer storage and formats them into the order required for the Memjet printhead. It is responsible for supplying the pre-calculated nozzleEnable bits to the Memjet interface for the eventual printing of the page.

The printing uses a double buffering scheme for preparing and accessing the dot-bit information. While one line is being loaded into the first buffer, the pre-loaded line in the second buffer is being read in Memjet dot order. Once the entire line has been transferred from the second buffer to the printhead via the Memjet interface, the reading and writing processes swap buffers. The first buffer is now read and the second buffer is loaded up with the new line of data. This is repeated throughout the printing process, as can be seen in the conceptual overview of Figure 51.

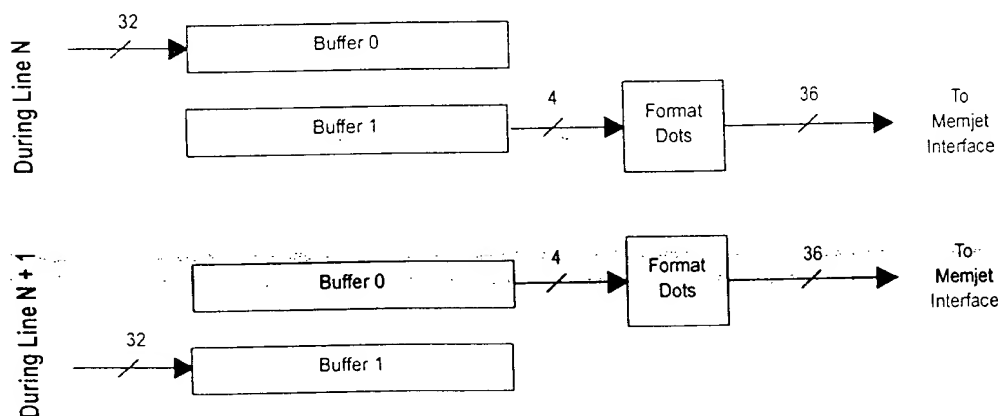


Figure 51. Conceptual Overview of Double Buffering During Print Lines N and N+1

The size of each buffer is 14KBytes to cater for the maximum line length of 18 inches in 4 colors ($18 \times 1600 \times 4$ bits = 115,200 bits = 14,400 bytes). The size for both Buffer 0 and Buffer 1 is 28.128 KBytes. While this design allows for a maximum print length of 18 inches, it is trivial to reduce the buffer size to target a specific application.

The actual implementation of the LLFU is shown in Figure 52. Since one buffer is being read from while the other is being written to, two sets of address lines must be used. The 32-bits *DataIn* from the common data bus are loaded depending on the *WriteEnables*, which are generated by the State Machine in response to the DMA Acknowledges.

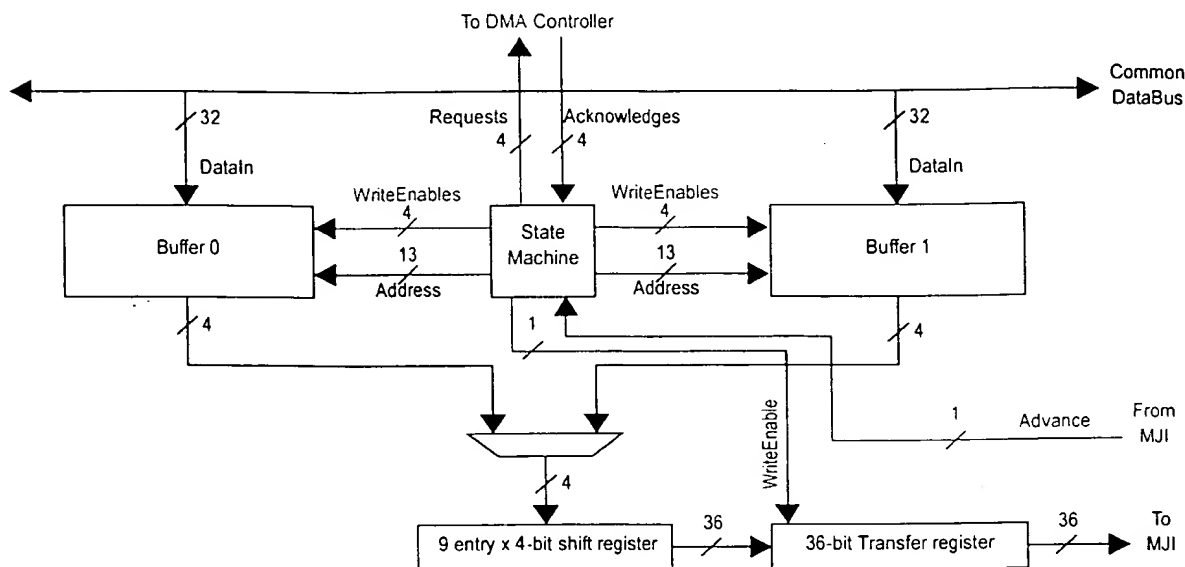


Figure 52. Structure of LLFU

A multiplexor chooses between the two 4-bit outputs of Buffer 0 and Buffer 1, and sends the result to a 9-entry by 4-bit shift register. After a maximum of 9 read cycles (the number depends on the number of segments written to per transfer), and whenever an Advance pulse comes from the MJ1, the current 36-bit value from the shift register is gated into the 36-bit Transfer register, where it can be used by the MJ1.

Note that not all the 36 bits are necessarily valid. The number of valid bits of 36 depends on the number of colors in the printhead, the number of segments, and the breakup of segment groups (if more than one segment group). For more information, see Section 14.2 on page 63.

A single line in an *L*-inch *C*-color printhead consists of **1600L C-color dots**. At 1 bit per colored dot, a single print-line consists of **1600LC bits**. The LLFU is capable of addressing a maximum line size of 18 inches in 4 colors, which equates to 108,800 bits (14 KBytes) per line. These bits must be supplied to the MJ1 in the correct order for being sent on to the printhead. See Section 14.2.1 on page 64 for more information concerning the Load Cycle dot loading order, but in summary, 2LC bits are transferred to the printhead in **SegmentGroups** transfers, with a maximum of 36 bits per transfer. Each transfer to a particular segment of the printhead must load all colors simultaneously.

15.4.1 Buffers

Each of the two buffers is broken into 4 sub-buffers, 1 per color. The size of each sub-buffer is 3600 bytes, enough to hold 18-inches of single color dots at 1600 dpi. The memory is accessed 32-bits at a time, so there are 900 addresses for each buffer (requiring 10 bits of address).

All the even dots are placed before the odd dots in each color's buffer, as shown in Figure 53. If there is any unused space it is placed at the end of each color's buffer.




Figure 53. Conceptual Structure of Buffer

The *amount of memory* actually used is directly related to the printhead length. If the printhead is 18 inches, there are 1800 bytes of even dots followed by 1800 bytes of odd dots, with no unused space. If the printhead is 12 inches, there are 1200 bytes of even dots followed by 1200 odd dots, and 1200 bytes unused.

The *number of sub-buffers* gainfully used is directly related to the number of colors in the printhead. This number is typically 3 or 4, although it is quite feasible for this system to be used in a 1 or 2 color system (with some small memory wastage). In a desktop printing environment, the number of colors would be 4: Color1=Cyan, Color2=Magenta, Color3=Yellow, Color4=Black.

The addressing decoding circuitry is such that in a given cycle, a single 32-bit access can be made to all 4 sub-buffers - either a read from all 4 or a write to one of the 4. Only one bit of the 32-bits read from each color buffer is selected, for a total of 4 output bits. The process is shown in Figure 54. 15 bits of address allow the reading of a particular bit by means of 10-bits of address being used to select 32 bits, and 5-bits of address choose 1-bit from those 32. Since all color buffers share this logic, a single 15-bit address gives a total of 4 bits out, one per color. Each buffer has its own WriteEnable line, to allow a single

32-bit value to be written to a particular color buffer in a given cycle. The 32-bits of DataIn are shared, since only one buffer will actually clock the data in.




Figure 54. Logical Structure of Buffer

Note that regardless of the number of colors in the printhead, 4 bits are produced in a given read-cycle (one bit from each color's buffer).

15.4.2 Address Generation

15.4.2.1 Reading

Address Generation for reading is straightforward. Each cycle we generate a bit address which is used to fetch 4 bits representing 1-bit per color for a particular segment. By adding 400 to the current bit address, we advance to the next segment's equivalent dot. We add 400 (not 800) since the odd and even dots are separated in the buffer. We do this firstly SegmentGroups sets of SegmentsPerXfer times to retrieve the data representing the even dots (the dot data is transferred to the MJ1 36 bits at a time) and another SegmentGroups sets of SegmentsPerXfer times to load the odd dots. This entire process is repeated 400 times, incrementing the start address each time. Thus all dot values are transferred in the order required by the printhead in $400 \times 2 \times \text{SegmentGroups} \times \text{SegmentsPerXfer}$ cycles.

In addition, we generate the TransferWriteEnable control signal. Since the LLFU starts before the MJ1, we must transfer the first value before the Advance pulse from the MJ1. We must also generate the next value in readiness for the first Advance pulse. The solution is to transfer the first value to the Transfer register after SegmentsPerXfer cycles, and then to stall SegmentsPerXfer-cycles later, waiting for the Advance pulse to start the next SegmentsPerXfer cycle group. Once the first Advance pulse arrives, the LLFU is synchronized to the MJ1. However, the LineSync pulse to start the next line must arrive at the MJ1 at least 2SegmentsPerXfer cycles after the LLFU so that the initial Transfer value is valid and the next 32-bit value is ready to be loaded into the Transfer register.

The read process is shown in the following pseudocode:

```

DoneFirst = FALSE
For DotInSegment0 = 0 to 400
  CurrAdr = DotInSegment0
  XfersRemaining = 2 x SegmentGroups
  DotCount = SegmentsPerXfer
  Do
    V1 = DotCount = 0
    TransferWriteEnable = (V1 AND NOT DoneFirst) OR ADVANCE
    Stall = V1 AND (NOT TransferWriteEnable)
    If (NOT Stall)
      Shift Register=Fetch 4-bits from CurrReadBuffer:CurrAdr
      CurrAdr = CurrAdr + 400
      If (V1)
        DotCount = SegmentsPerXfer - 1
        XfersRemaining = XfersRemaining - 1
      Else
        DotCount = DotCount - 1
      EndIf
    EndIf
  Until (XfersRemaining=0) AND (NOT Stall)
EndFor

```

The final transfer may not be fully utilized. This occurs when the number of segments per transfer does not divide evenly into the actual number of segments in the printhead. An example of this is the 8½" printhead, which has 17 segments. Transferring 9 segments each time means that only 8 of the last 9 segments will be valid. Nonetheless, the timing requires the entire 9th segment value to be generated (even though it is not used). The actual address is therefore a don't care state since the data is not used.

Once the line has finished, the CurrReadBuffer value must be toggled by the processor.

15.4.2.2 Writing

The write process is also straightforward. 4 DMA request lines are output to the DMA controller. As requests are satisfied by the return DMA Acknowledge lines, the appropriate 8-bit destination address is selected (the lower 5 bits of the 15-bit output address are *don't care* values) and the acknowledge signal is passed to the correct buffer's WriteEnable control line (the Current Write Buffer is -CurrentReadBuffer). The 10-bit destination address is selected from the 4 current addresses, one address per color. As DMA requests are satisfied the appropriate destination address is incremented, and the corresponding TransfersRemaining counter is decremented. The DMA request line is only set when the number of transfers remaining for that color is non-zero.

The following pseudocode illustrates the Write process:

```

CurrentAdr[1-4] = 0
While (ColorXfersRemaining[1-4] are non-zero)
  DMARequest[1-4] = ColorXfersRemaining[1-4] NOT = 0
  If DMAAcknowledge[N]
    CurrWriteBuffer:CurrentAdr[N] = Fetch 32-bits from data bus
    CurrentAdr[N] = CurrentAdr[N] + 1
    ColorXfersRemaining[N] = ColorXfersRemaining[N] - 1 (floor 0)
  EndIf
EndWhile

```

15.4.3 Registers

The following interface registers are contained in the LLFU:

Table 21. Line Load/Format Unit Registers

Register Name	Description
SegmentsPerXfer	The number of segments whose dots must be loaded before each transfer. This has a maximum value of 9.
SegmentGroups	The number of segment groups in the printhead. This has a maximum number of 4.
CurrentReadBuffer	The current buffer being read from. When Buffer0 is being read from, Buffer1 is written to and vice versa. Should be toggled with each AdvanceLine pulse from the MJ1.
Go	Bits 0 and 1 control the starting of the read and write processes respectively. A non-zero write to the appropriate bit starts the process.
Stop	Bits 0 and 1 control the stopping of the read and write processes respectively. A non-zero write to the appropriate bit stops the process.
Stall	This read-only status bit comes from the LLFU's Stall flag. The Stall bit is valid when the write Go bit is set. A Stall value of 1 means that the LLFU is waiting for the ADVANCE pulse from the MJ1 to continue. The CPU can safely start the LSGU for the first line once the Stall bit is set.
ColorXfersRemaining[1-4]	The number of 32-bit transfers remaining to be read into the specific Color[N] buffer.

15.5 CONTROLLING A PRINT

When controlling a print the CPU programs and starts the LLFU in read mode to ensure that the first line of the page is transferred to the buffer. When the interrupts arrive from the DMA controller, the CPU can switch LLFU buffers, and program the MJ1. The CPU then starts the LLFU in read/write mode and starts the MJ1. The CPU should then wait a sufficient period of time to ensure that other connected printer controllers have also started their LLFUs and MJIs (if there are no other connected printer controllers, the CPU must wait until the Stall bit of the LLFU is set, a duration of 2SegmentsPerXfer cycles). The CPU can then program the LGSU to start the synchronized print. As interrupts arrive from the DMA controllers, the CPU can reprogram the DMA channels, swap LLFU buffers, and restart the LLFU in read/write mode. Once the LLFU has effectively filled its pipeline, it will stall until the next Advance pulse from the MJ1. The MJ1 does not have to be touched during the print.

If for some reason the CPU wants to make any changes to the MJ1 or LLFU registers during an inter-line period it should ensure that the current line has finished printing/loading by polling the status bits of the MJ1 and the Go bits of the LLFU.

REFERENCES

16 Silverbrook References

- [1] Silverbrook Research, *Authentication Chip*, 1998
- [2] Silverbrook Research, *Authentication of Consumables*, 1998
- [3] Silverbrook Research, *CePrint Product Design Description*, 1998
- [4] Silverbrook Research, *Memjet*, 1998
- [5] Silverbrook Research, *Netpage Pen Architecture*, 1999

17 Other References

- [6] ANSI/EIA 538-1988, *Facsimile Coding Schemes and Coding Control Functions for Group 4 Facsimile Equipment*, August 1988
- [7] Farrell, J., "How to Allocate Bits to Optimize Photographic Image Quality", Proceedings of IS&T International Conference on Digital Printing Technologies, 1998, pp.572-576
- [8] Humphreys, G.W., and V. Bruce, *Visual Cognition*, Lawrence Erlbaum Associates, 1989, p.15
- [9] ISO/IEC 19018-1:1994, *Information technology - Digital compression and coding of continuous-tone still images: Requirements and guidelines*, 1994
- [10] Loeffler, C., A. Ligtenberg and G. Moschytz, "Practical Fast 1-D DCT Algorithms with 11 Multiplications", Proceedings of the International Conference on Acoustics, Speech, and Signal Processing 1989 (ICASSP '89), pp.988-991
- [11] Thompson, H.S., *Multilingual Corpus 1* CD-ROM, European Corpus Initiative
- [12] Urban, S.J., "Review of standards for electronic imaging for facsimile systems", *Journal of Electronic Imaging*, Vol.1(1), January 1992, pp.5-21
- [13] Wallace, G.K., "The JPEG Still Picture Compression Standard", *Communications of the ACM*, Vol.34, No.4, April 1991, pp.30-44

Netpage Pen

Design Description

draft version 0.3, 7 September 1999



Silverbrook Research Pty Ltd
393 Darling Street, Balmain
NSW 2041 Australia
Phone: +61 2 9818 6633
Fax: +61 2 9818 6711
Email: info@silverbrook.com.au

Confidential

Contents

PEN CONTROLLER	1
1 Pen Controller Architecture	2
NETPAGE ENCODING AND DECODING	3
2 Netpage Encoding.....	4
2.1 Position Tag Structure.....	6
2.2 Position Tag Encoding	9
3 In-Pen Processing.....	12
3.1 Position Tag Decoding	12
3.2 Data Encryption	24
REFERENCES	26
4 Silverbrook References.....	27
5 Other References	28

PEN CONTROLLER

1 Pen Controller Architecture

The Netpage Pen controller consists of a controlling processor, a radio transceiver, a tilt sensor, a nib pressure sensor, a IR image sensor, flash memory, and 512KB of DRAM, as illustrated in Figure 1.

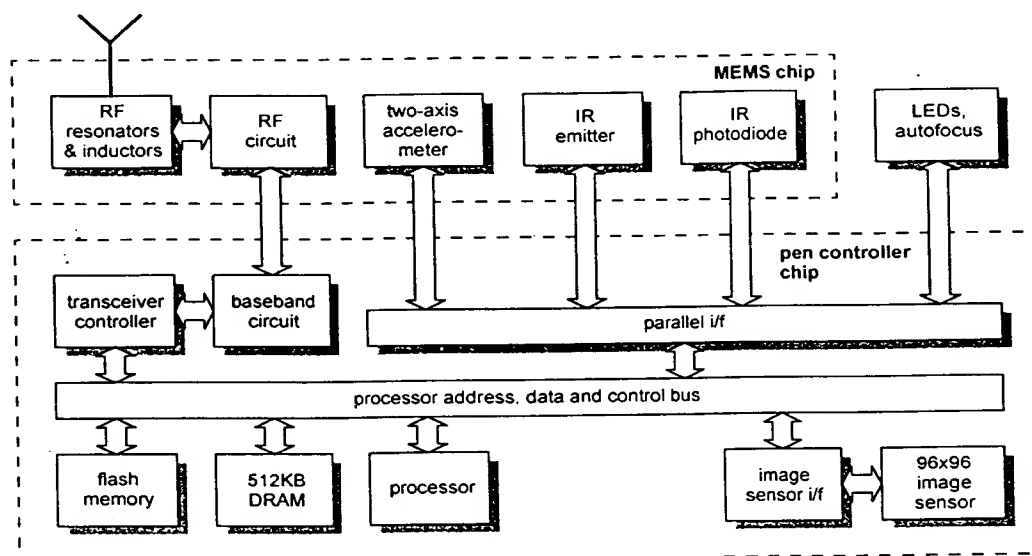


Figure 1. Pen controller architecture

The controlling processor captures and decodes IR position tags from the page via the image sensor, monitors the tilt and pressure sensors, controls the autofocus voice-coil, controls the user interface (tricolor-LED), and handles wireless communication with the local Netpage Printer. It consists of a medium-performance (~25MHz RISC) general-purpose microprocessor.

Two-axis tilt sensing is provided by a two-axis accelerometer. Nib pressure sensing is provided by an IR emitter and photodiode pair in conjunction with a reflector coupled with the sprung nib.

The processor, digital transceiver components, 96x96 image sensor, flash memory and 512KB DRAM are integrated in a single ASIC. The MEMS and analog RF components, accelerometers, and the IR emitter/photodiode are integrated in a single MEMS chip, also used in the Netpage Printer.

The radio transceiver communicates in the unlicensed 900MHz band normally used by cordless telephones, and uses frequency hopping and collision detection to provide interference-free communication.

NETPAGE ENCODING AND DECODING

2 Netpage Encoding

Hidden information on a printed Netpage encodes the unique id of the page instance, together with position information which tells the pen which part of the page the user is interacting with. The Netpage Network can correlate any reference to any part of any page ever printed with the contents of the page, solely on the basis of the page id and the position.

The entire page is tiled with position tags. Each position tag must be independent, to allow for single-point virtual button and hyperlink presses, and so contains both the local x-y position and the global page id. In addition particular flags are associated with each tag, such as whether the pen should turn on its light while over the tag.

The tag data (page id, a local x-y position, and control flags) is error-correctably encoded as a square binary matrix which is printed as an array of bi-level macrodots. Each macrodot itself consists of a square of one or more printed dots to ease image processing.

The position encoding supports a particular page size. Where the physical page size exceeds this virtual page size, the physical page is simply tiled with multiple virtual pages, each with a different page id. Figure 2 illustrates the concept of a page tiled with position tags. The contents of each tag are not shown on this illustration.




Figure 2. Part of a page tiled with position tags

The tags are printed with an infrared-absorptive ink that can be read by the pen device. Since black ink is IR absorptive, limited functionality can be provided on offset-printed pages using black ink on otherwise blank areas of the page - for example to encode buttons. Alternatively an invisible infrared ink can be used to print the position tags over the top of a regular page. However, if invisible IR ink is used, care must be taken to ensure that any other printed information on the page is printed in infrared-transparent CMY ink, for black ink will obscure the infrared tags. The monochromatic scheme was chosen to maximize dynamic range in blurry reading environments.

Simple virtual button presses don't specifically require invisible ink. Buttons could be visibly coded since the real estate is available. Hyperlinks require invisible ink, but impose no other requirements on the tag encoding.

Table 1. Encoding parameters

parameter	definition	value
p	virtual page width (inches)	12 ^a
d	dot pitch (dots per inch)	1600
k	macrodot width (dots)	4
b	registration border width (macrodots)	2
r	data redundancy factor	1.5
m	coordinate precision (bits)	-
g	page id precision (bits)	-
s	tag width (macrodots)	-
n	tags per page ^b	-

a. US Letter page length

b. per dimension

Given the encoding parameters defined in Table 1, the tag width is given by:

$$s = b + \sqrt{r(2m + g)}$$

Given a particular tag width, the number of tags per page in each dimension is given by:

$$n = \frac{pd}{ks}$$

This in turn yields the required coordinate precision:

$$m = \log_2 n$$

If 10 billion people all generate 100,000 pages per annum for 1000 years, they will generate 10^{18} pages, or approximately 2^{60} pages. A page id precision of 64 bits should therefore be sufficient, notwithstanding issues of efficient contiguous page id allocation.

Assuming a page id precision of 64 bits, and other parameter values as given in Table 1, the equations converge on a tag width of 17 macrodots, 283 tags per page dimension, and a coordinate precision of 9 bits. This precision supports a maximum page dimension of 21.7" or 552mm.

Given a particular tag width, the position resolution, *in millimeters*, is given by:

$$\frac{25.4ks}{d}$$

and, *in points*, is given by:

$$\frac{72ks}{d}$$

This yields a position resolution of 3 points (1.08mm at 1600 dpi).

Magazine-quality printed text normally has a size of 10 points. Forms filled in by hand normally allow for handwritten text with a size of about 20 points. A position resolution of 3 points therefore translates to a character-oriented resolution of 3.3 for printed text and 6.6 for handwriting recognition. This position resolution is increased to the 200 dpi required for handwriting recognition by taking into account the position of the tag within the captured image area.

With a position tag width of 17 macrodots, and 4x4 dots (at 1600 dpi) per macrodot, the raw position tag gives about 24 positions per inch. Consequently any positioning scheme must achieve 10 times better resolution based on tag position within the sensed image.

2.1 POSITION TAG STRUCTURE

Each position tag encodes the following information:

Table 2. Data Encoded in a Position Tag

Name	#Bits	Description
Pageld	64	Defines the Netpage page instance id
x	9	Defines the x coordinate within the given Pageld
y	9	Defines the y coordinate within the given Pageld
ActiveArea	1	Defines whether the pen should turn its light on while over the tag
PressureSensitiveArea	1	Determines whether the pen should return continuous pressure or not while over the tag. 0 = don't return pressure readings. 1 = return pressure readings
Reserved	6	For future use. Store as 0 for this version.
TOTAL	90	

Each position tag is a 17x17 array of macrodots, with each macrodot being a 4x4 square of 1600 dpi dots, as shown Figure 3.




Figure 3. Macrodots in a Position Tag

Since the dots are monochrome IR-absorptive dots, printed on an IR-reflective background, a "black macrodot" is physically different from a "white macrodot". Figure 4 illustrates a magnified view of macrodots.




Figure 4. Magnified Macrodots

The position tag structure is based on concepts used in the public-domain Aztec 2D Barcode [7], invented by Andy Longacre of Welch Allyn Inc. in 1995. While the position tag uses the same bull's-eye structure mechanism from the Small Aztec Symbol for symbol location and orientation, we define our own interpretation for the mode bits and do not allow variable length structures. In addition, the position tag is a 17×17 array, which is not possible in a pure Aztec code.

Figure 5 shows the high-level structure of a position tag in relation to its 17×17 macrodot array. The target/orientation area consists of a 5-layer bull's-eye with 3-bit orientation markers on the outer corners. The format of this area is identical to that used in the Small Aztec Barcode although it is placed off-center. The Mode Data Area is the area between the orientation markers in what would otherwise be the 6th layer of the bull's-eye target. The remaining area is used to hold the actual data for the position tag.

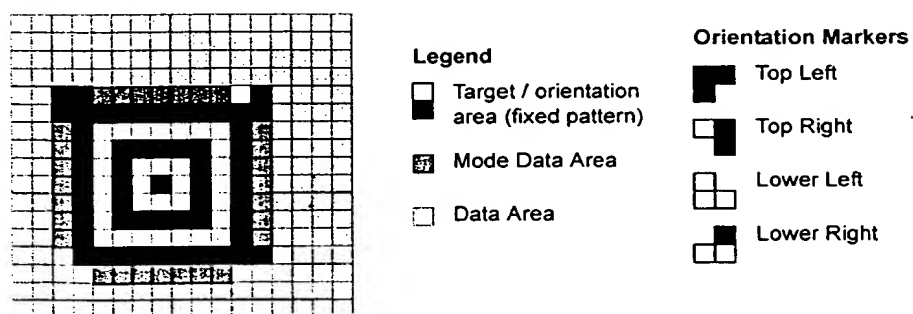


Figure 5. High Level Structure of a Position Tag

2.1.1 Target and Orientation Area

The target and orientation area consists of the fixed pattern shown in black and white in Figure 5. The format of this area is identical to that used in the Small Aztec Barcode.

The black positions in are black macrodots, and the white positions are white macrodots. The 5-layer bull's-eye pattern is easily found in a 2D image by scans for topological connectivity, and is then useful for pinpointing the exact center regardless of orientation, and for determining the main axes and local x-y dimensions. On the four corners of the target are 3-bit clusters of orientation bits, which allow the tag's orientation to be quickly determined.

2.1.2 Mode Data Area

The bits in the layer immediately adjoining the target (other than the orientation bits) comprise a 28-bit string, starting upper left and circling clockwise. This area is known as the Mode Data Area. The relative position of the 28 Mode Data bits can be seen in Figure 6.




Figure 6. Mode Data Area

The 28-bit string encodes meta-information used to decode the remainder of the position tag. Of particular note is the first bit of the meta information. If this bit is 0, then the decoding of the data area is as described here. Setting this bit to 1 allows for redefinition of the position tags at some later date while remaining compatible with existing pages. The meta information and its encoding to 28 bits is described in Section 2.2.1.

2.1.3 Data Area

When the first bit of the meta information decoded from the Mode Data Area is 0, then the data area is defined to be the 168 bits as shown in Figure 7.

The encoded message data is placed onto the tag in 2-cell high layers, starting at the upper left and spiraling clockwise out to the edge of the symbol. The data is therefore interpreted as a sequence of related “dominos”, each 1 wide and 2 tall, with their more significant bit always further from the target area (a black macrodot represents a “1” bit, and a white macrodot represents a “0” bit). Figure 7 illustrates the sequence and orientation of the dominos when turning the corners in the data area.




Figure 7. Domino Mapping and Sequencing of Data Area

The systematic domino-based layout of the data simplifies both encoding and decoding at the graphical level. The encoding mechanism for the data is described in Section 2.2.2, although it is worth noting here that decoding errors and tag damage are expected to occur more towards the edges of the tag.

Note that the 17×17 structure of the position tag means that with reference to the strict Aztec Code, the first layer of data is complete, but the second layer of data is only half full. If the second layer was completely full the position tag would be a 19×19 square cell with the target area in the center.

2.2 POSITION TAG ENCODING

2.2.1 The Mode Message

The 28-bit string in the Mode Area is used to hold 8 mode bits. These 8 mode bits describe how the remainder of the position tag bits are encoded.

The 8 mode bits are encoded into 28 bits via Reed-Solomon encoding. The Mode Area therefore holds the 8 mode bits and 20 additional check bits. The 8 mode bits are parsed into two 4-bit words, and then 5 additional check words are computed by systematic Reed-Solomon encoding over the Galois field GF(16) based on a prime modulus polynomial of: $x^4 + x + 1$ (=19 decimal). The generator polynomial of $(x-2^1) \dots (x-2^4)$ is:

$$x^4 + 11x^3 + 4x^2 + 6x + 2$$

Reed-Solomon encoding is chosen for its ability to deal with burst errors and effectively detect and correct errors using a minimum of redundancy. Reed-Solomon coding is discussed in detail in [8], [4], and [3]. The reader is advised to refer to these sources for background information.

Table 3 shows the interpretation of the mode bits. Note that the first bit of the mode data is a version bit, and determines the interpretation of the remainder of the mode bits and the tag data. The Version bit should therefore be set to 0. Later interpretations of the tag can be defined at a later date via the version bit.

Table 3. Interpretation of Mode Bits

#Bits	Name	Description
0	Version	0 = the interpretation of the remaining 7 bits and the data area is as described in this document 1 = reserved for future use
1-2	NumMsgWords	Defines how many 6-bit codewords of the data area are message codewords, and how many are check-word codewords. This number is determined by the encoding scheme used for data (see Section 2.2.2). Values are: 00 = 15 data message codewords 01 = 16 data message codewords 10 = 17 data message codewords 11 = 18 data message codewords
3-7	Reserved	0

2.2.2 The Data Message

The Data Area contains 168 raw bits, which is an encoded form of the 90 bits of position tag data (see Table 2). The 168 bits are represented by 104 bits in the first layer surrounding the target area and 64 bits in the half layer surrounding the first layer.

The 90 bits of position tag data are encoded into 168 bits in two steps. The first step breaks the 90 bits into codewords of 6 bits such that there are no codewords with all 0s or all 1s.

The second step Reed-Solomon encodes the data over the Galois field GF(64), which generates additional 6-bit check words.

2.2.2.1 Step 1: Message Encoding

In the first step of message encoding, the 90 data bits are arranged into a sequence of 6-bit message codewords in a generally direct fashion, starting at the most significant bit of the first codeword, with two key exceptions:

- whenever the first 5 bits placed in a codeword are all “0”s, then a dummy “1” is inserted into that codeword’s LSB and the following message bit starts off the next codeword.
- whenever the first 5 bits placed in a codeword are all “1”s, then a dummy “0” is inserted into that codeword’s LSB and the following message bit starts off the next codeword.

In the end, the character and byte boundaries in the original 90-bit message have no necessary relationship with the codeword boundaries. Up to 5 bits may remain unfilled in the final message codeword, and they are padded out with “1”s (and possibly a final dummy “0” if necessary) to eliminate any ambiguity.

The code-forming rules are designed to never create a *message* codeword of all 0s or all 1s, but the error encoding adds on additional codewords of any value. Thus during decoding, an occurrence of those illegal values within the message region (but *not* within the check region) can be regarded as a correctable *erasure* (this is more useful than not knowing which codeword is in error).

In the case of our 90-bit original messages, there are two cases to consider:

- the *best* case is that 15 codewords are generated for the data message
- the *worst* case is that 18 codewords are generated for the data message

The codeword based message therefore consumes D 6-bit codewords, where D is in the range 15-18.

2.2.2.2 Step 2: Generate Check Words

The second step of message encoding involves generating the check words for the data message. Since the data message consumes between 15 and 18 6-bit codewords of the 28 available codewords, the remaining 13 to 10 codewords can be used for check words.

The K additional check words are computed by systematic Reed-Solomon encoding over the Galois field GF(64) based on a prime modulus polynomial of: $x^6 + x + 1$ (= 67 decimal). The generator polynomial is simply expanded out to $(x-2^1) \dots (x-2^K)$ as needed at the time of printing. The check words are not adjusted in the same way as the message bits are adjusted.

2.2.2.3 Step 3: Place into Data Area

The final message to be stored in the data area of a position tag consists of 28 6-bit codewords made up of D message codewords followed by K check codewords. The exact number of message codewords, D , is stored into the 2-bit NumMsgWords component of the Mode information for the position tag.

The resulting sequence of codewords is parsed anew into a sequence of 2-bit dominos (see Figure 8) and is then taken in reverse order and graphically arranged spiralling clockwise and outwards through the data area (see Figure 7). As decoding errors and tag damage are expected to occur more towards the edges of the tag, by reversing the codeword sequence, the message codewords occupy those edge regions.




Figure 8. Relationship of 6-bit codeword to generated dominos

One macrodot is placed into the data area per bit. The placement of an IR-absorptive ink macrodot represents a "1". If there is no IR-absorptive ink placed on the page at the particular location, then that macrodot represents a "0".

3 In-Pen Processing

The pen must send the printer 100 samples per second. The mechanism for reading positional information from the position tags must work at this 100 Hz rate. The following tasks must be performed each sample:

- capture time, tilt and pressure
- capture the image
- locate the position tag
- decode the position tag
- build stroke information
- activate the light as required
- encrypt the stroke data
- transmit the encrypted data

When a pen-down is registered, we capture the first three frames at the two extreme focus values and the middle focus value. After that we perform regular auto-focus based on a sharpness metric. In this way we are more likely to capture a "single-click" event.

While the pen is in range of the printer, partially completed strokes are transmitted as they are being formed. When the pen moves out of range, stroke information is buffered within the pen (approximately 12 minutes of pen motion while physically on the page), and transmitted later.

Each pen contains a characteristics block which can be read by the printer. It contains at least the information shown in Table 4.

Table 4. Pen characteristics block

Name	Description
PenId	A unique pen identifier - unique across all pens
ModelId	The model number of this pen ^a
ManufacturerId	The manufacturer of this pen
TiltInfo	The relationship between tilt and pen position.

- a. This can be combined with ManufacturerId to allow the printer to download a set of pen characteristics. However it may be more convenient to store the characteristics (such as TiltInfo) in the pen.

3.1 POSITION TAG DECODING

3.1.1 Capture Time, Tilt and Pressure

The pen contains sensors which allow pen tilt and nib pressure to be determined.

3.1.1.1 Time

All measurements are made in the context of a particular timestamp. The pen contains a timer to allow equally time-spaced measurements to be made. Although we capture samples at a rate of 100 per second, we report time at millisecond resolution to allow for future improvements to the pen. Time resolution is provided at 32 bits.

3.1.1.2 Tilt

The pen contains two orthogonal passive accelerometers in the x-y plane perpendicular to the pen axis. These respond to gravity and allow two-axis tilt to be computed to an accuracy of at least 5 bits.

Tilt is required to determine the nib position. The printer uses the two tilt values and the pen's non-corrected pen position to determine the actual pen nib position. The Pen Characteristics Block contains the relationship required for the calculation. Offloading the correction calculation to the printer reduces pen complexity and price.

3.1.1.3 Pressure

The pen contains a pressure sensor which measures to an accuracy of at least 5 bits. From the pen's point of view, the pressure value is only used to determine whether the pen is on the paper or not. The printer is responsible for all other interpretations of pressure, such as determining whether an item on a page is being selected. The pen merely passes the pressure value on to the printer.

Each position tag contains a data bit called PressureSensitiveArea. If a given pen stroke includes position tags that have their PressureSensitiveArea bit set, then the pressure value is passed as part of the stroke information. If the PressureSensitiveArea is clear, then pressure is implied from a stroke in terms of pen-down and pen-up.

3.1.2 Capture the Image

The image needs to be captured for the pen position to be determined. This only needs to happen if the pen is on the page.

To capture an image, data must be transferred from the image sensor to a fixed image buffer. This process involves a single read and write per pixel. The minimum and maximum values encountered during the transfer are kept in variables MinPixelEncountered and MaxPixelEncountered, respectively, for later use in the data recovery process (see Section 3.1.4).

There are a number of general considerations that are part of the assumptions for reading in and decoding a position tag from an image sensor.

With regards to the image sensor itself, there are two calculations to consider: the size of the sensed area, and the resolution of the image sensor. As described in Section 3.1.2.1 and Section 3.1.2.2, the sensed area has a required size of 193×193 dots (or 3.06mm × 3.06mm), and the image sensor has a required resolution of 96×96 pixels.

Furthermore, as discussed in Section 3.1.2.3, the sensed image may be blurry, and the sense of ambiguous macrodots may have to be inferred.

A 96×96 image contains 9216 8-bit pixels. We round this up to 10,000 pixels for calculation purposes. The transfer process involves on average 6.5 cycles per pixel for a total of 65,000 cycles per position tag. On a 25 MHz processor this represents 26% of the available bandwidth.

3.1.2.1 Sensed Area Size

The basic position tag consists of 17×17 macrodots. Since each macrodot is a 4×4 array of printed dots, each position tag in fact consists of 68×68 printed dots.

To ensure that one entire tag is always within the window of the image sensor, the image sensor window must be the width of two tags. This requirement increases the sensed dot area to 136×136 .

Finally, since the pen must detect positional tags at any rotation, we must consider the worst case rotation of 45 degrees. This requirement increases the sensed dot area to 193×193 .

At a printed dot size of 1600 dpi, the image sensor must sense a printed area containing 193×193 dots, which equates to 3.06mm (0.12 inches) in each dimension.

3.1.2.2 Image Sensor Resolution

The basic positional tag is 17×17 macrodots. This equates to a basic sensor resolution of 17×17 pixels.

To ensure that one entire tag is always within the window of the image sensor, the image sensor window must be the width of two tags. This requirement increases the sensor resolution to 34×34 .

Since the pen must detect positional tags at any rotation, we must consider the worst case rotation of 45 degrees. This requirement increases the sensor resolution to 48×48 .

Finally, to satisfy Nyquist's Theorem with respect to macrodots, we must oversample at least at twice the macrodot resolution. This leads to a sensor resolution of 96×96 .

Each sampled pixel is 1 byte (8 bits). We must also assume the least significant 2 bits of each pixel can contain noise. Decoding algorithms must therefore be noise-tolerant.

3.1.2.3 Blurry Image

The pen optics provide reasonable depth-of-field, and the pen includes an auto-focus mechanism to handle changes in pen-paper separation due to varying tilt etc., but several factors may still contribute to blurriness in the captured image, including tardy auto-focus response, and varying pen-paper separation due to extreme tilt or warped paper.




Figure 9. Effect of different degrees of blurring

Blurring can prevent definitive macrodot recognition. Figure 9 shows the effect of blurring on a position tag composed of macrodots, ranging from no blurring to severe blurring. Note how isolated macrodots fade as the blurring becomes more severe.

As the blurring increases, the more a given macrodot is influenced by the surrounding macrodots. Consequently the dynamic range for a particular macrodot decreases. Consider a "white" macrodot and a "black" macrodot, each surrounded by all possible sets of macrodots. The 9 macrodots are blurred, and the center macrodot sampled. Figure 10 shows the distribution of resultant center macrodot pixel values for black and white macrodots.




Figure 10. Black-white macrodot ambiguity due to blurring

The curve in Figure 10 is computed for an illustrative degree of blurring. The curve from pixel values 0-180 shows the range of black macrodots. The curve from pixel values 75-250 shows the range of white macrodots. The greater the degree of blurring, the more the two curves shift towards the center of the range and therefore the greater the intersection area, which means the more difficult it is to determine whether a given macrodot is black or white. A pixel value at the center point of intersection is ambiguous - the macrodot is equally likely to be black or white.




Figure 11. Relationship between data non-recovery and codeword error rate

As the blurring increases, the likelihood of a read bit error increases. Fortunately, the Reed-Solomon decoding algorithm can cope with bit errors gracefully up to t symbol errors. Figure 11 shows the predicted number of Reed-Solomon codewords that cannot be recovered given a particular symbol error rate. Notice how the Reed-Solomon decoding

scheme performs well and then substantially degrades. If there is no Reed-Solomon block duplication, then only 1 codeword needs to be in error for the data to be unrecoverable. Of course, with the inclusion of erasure detection (via invalid codewords), the chance of correctly decoding the data increases.

Figure 11 only links codeword errors to the total number of Reed-Solomon codewords in error. There is a trade-off between the amount of blurring that can be corrected, and the number of genuine errors due to dirt and paper damage that can be corrected. Since all error detection and correction is performed by the Reed-Solomon decoder, there are a finite number of errors per Reed-Solomon data block that can be corrected. The larger the number of errors introduced through blurring, the smaller the number of genuine errors that can be corrected.

3.1.3 Locate the Position Tag

The captured image is a 96×96 pixel map of 8-bit samples. Before any position information can be decoded, the position tag must be located within the image. Note that the tag can have any orientation.

The first step in locating the tag is to locate a bull's-eye target (Figure 5) within the image. Since we are imaging a wider area than a single tag it is possible that two targets lie within the image area.




Figure 12. Area of interest for target location

Since an entire tag must be completely contained within the image, the data area outside the target must be visible. As a result, the area of interest for locating targets is actually a subset of the whole image. Instead of searching a 96×96 image, we only need search an 88×88 image (a reduction of 15% from 9216 samples to 7744).

The 5-layer bull's-eye pattern of the target is easily found in a 2D image by scans for topological connectivity, and is then useful for pinpointing the exact center regardless of orientation, and for determining the main axes and local x-y dimensions.

3.1.3.1 Scan for Topological Connectivity

Prior to image traversal we initialize an 88×88 array of flags with 0. This array keeps the visited status of the corresponding pixel in the image. If the flag is 0 we have not yet visited the pixel. If it is 1 or 2, then we have visited the pixel and bit 1 of the flag contains the sense of the bit (1 for black, 0 for white).

We then start with the top left pixel and compare adjacent pixels in order to build a connectivity structure for the top left pixel. If the top left pixel is considered to be white, then the connectivity structure will represent all connected pixels that are also white. If the top pixel is considered to be black, then the connectivity structure will represent all connected pixels that are also black. As edges are met due to differences in color sense, adjacent connectivity structures are built.

The order of pixel traversal in the image simply follows a standard pixel fill algorithm. There are many such algorithms to be found since this is a useful operation for paint programs. Near-optimal algorithms such as [1] and [2] only read each pixel once on average.

In our case we must also consider the notion of color sense. Target location does not require the sophistication of the macrodot reading algorithm for mode bits and data bits for determining if a macrodot is black or white (Section 3.1.4.2 and Section 3.1.4.4). The minimum and maximum pixels encountered during the image capture were stored in `MinPixelEncountered` and `MaxPixelEncountered` respectively, and these are used to set the threshold for black/white as $\text{MidRange} = (\text{MinPixelEncountered} + \text{MaxPixelEncountered})/2$. Anything greater than or equal to `MidRange` is considered to be white, and anything less than `MidRange` is considered to be black.

As each connectivity structures is completed, it is compared against the expected values for the 5 bull's-eye layers surrounding the target center. Possible matches are added to one list, while non-matches are discarded.

The total time taken for processing the image is approximately 6 cycles per sample pixel. This equates to 46,464 cycles ($88 \times 88 \times 6$), or 19% of the available processor bandwidth.

3.1.3.2 Determine Target Axes

The connectivity structures can now be examined together in order to locate the target. The identification of two or more concentric structures is enough to determine the x and y axes.

Taking one of the completely isolated concentric structures, 2 points on each of the 4 sides are chosen where the distance separating them is at least 3 pixels. For each point, the pixel of maximum whiteness or maximum blackness is chosen as the estimate of the center of the line. The estimate should be within 1 pixel of the actual center.

The process of building a more accurate position for the line center involves reconstructing the continuous signal for 7 scanline slices of the line, 3 to either side of the estimated center. The 7 maximum values found (one for each of these pixel dimension slices) are then used to reconstruct a continuous signal in the column dimension and thus to locate the maximum value in that dimension.

```
// Given estimates column and pixel, determine a
// betterColumn and betterPixel as the center of the target
for (y=0; y<7; y++)
{
    for (x=0; x<7; x++)
        samples[x] = GetPixel(column-3+y, pixel-3+x)
    FindMax(samples, pos, maxVal)
    reSamples[y] = maxVal
    if (y == 3)
        betterPixel = pos + pixel
}
```

```
FindMax(reSamples, pos, maxVal)
betterColumn = pos + column
```

FindMax is a function that reconstructs the original one-dimensional signal based sample points and returns the position of the maximum as well as the maximum value found. The method of signal reconstruction/resampling used is the Lanczos3 windowed sinc function (see pages 157-158 in [6]). The function and kernel are shown in Figure 13:




Figure 13. Lanczos3 windowed sinc function

The Lanczos3 windowed sinc function takes 7 (pixel) samples from the dimension being reconstructed, centered around the estimated position x , i.e. at $x-3$, $x-2$, $x-1$, x , $x+1$, $x+2$, $x+3$. We reconstruct points from $x-1$ to $x+1$, each at an interval of 0.1, and determine which point is the maximum. The position that is the maximum value becomes the new center. Due to the nature of the kernel, we only require 6 entries in the convolution kernel for points between x and $x+1$. We use 6 points for $x-1$ to x , and 6 points for x to $x+1$, requiring 7 points overall in order to get pixel values from $x-1$ to $x+1$ since some of the pixels required are the same.

Two accurate positions on each of 2 sides of the isolated bull's-eye layer give the x and y axes directly. Using 4 sides allows an average line to be calculated (simply the average of the two lines built from opposite sides).

The distance from the center of one bull's-eye layer macrodot line to the center of the next bull's-eye layer's macrodot line gives the distance between macrodots in terms of a Δx and a Δy with respect to the current orientation of the image.

3.1.3.3 Determine Target Center

In order to locate the actual center of the target there are two cases to consider.

- The center of the target is one of the connectivity structures
- The center of the target is *not* one of the connectivity structures

In the first case, we choose the blackest pixel of the estimated target center, which should be within 1 pixel of the actual center. The process of building a more accurate position for the target center involves reconstructing the continuous signal for 7 scanline slices of the target, 3 to either side of the estimated target center. The 7 maximum values found (one for each of these pixel dimension slices) are then used to reconstruct a continuous signal in the column dimension and thus to locate the maximum value in that dimension. The process is the same as described in Section 3.1.3.2 and uses the same filter as shown in Figure 13.

If there is no current estimation for the center, the center is considered to be the calculated center of the four lines used in the calculation of the axes in Section 3.1.3.2. This process can be repeated for another isolated bull's-eye layer and the two estimates for target center averaged. If the estimated center pixel is black, a more accurate center can potentially be obtained by following the procedure defined for the first case. However care must be taken since this pixel region is not completely isolated (or there would have been a target estimate in the connectivity structures). It may be enough to increase the accuracy in only one dimension.

3.1.4 Decode the Position Tag

With the position of the tag known, the task of decoding it can begin. Decoding a positional tag consists of 3 essential steps:

- Determining the orientation of the tag
- Extracting the tag's mode data and decoding it
- Extracting and decoding the data portion of the tag based on the mode data

3.1.4.1 Determine Orientation

On the four corners of the target are 3-bit clusters of orientation bits, which allow the tag's orientation to be quickly determined.

Using the Δx and Δy values (obtained in Section 3.1.3.2), and applying them to the accurate target center coordinate (obtained in Section 3.1.3.3), we are able to calculate the estimated positions of the 12 orientation macrodots.

Given the macrodot coordinate (fixed point) we sample 4 image pixels to arrive at a center pixel value via bilinear interpolation.

Once the center pixel value has been determined, we try to determine the bit value for that macrodot. To do so, we take the pixel values representing the macrodot centers to either side of the macrodot whose bit value is being determined, and attempt to intelligently guess the value of that center macrodot's bit value. Looking at the generalized blurring curve again (reproduced in Figure 14 from Figure 10), but with BlackMax and WhiteMin shown, there are three common cases to consider:




Figure 14. Black-white macrodot ambiguity due to blurring

- The macrodot's center pixel value is lower than WhiteMin, and is therefore definitely a black macrodot. The bit value is therefore definitely 1.

- The macrodot's center pixel value is higher than `BlackMax`, and is therefore definitely a white macrodot. The bit value is therefore definitely 0.
- The macrodot's center pixel value is somewhere between `BlackMax` and `WhiteMin`. The macrodot may be black, and it may be white. The value for the bit is therefore in question. A number of schemes can be devised to make a reasonable guess as to the value of the bit. These schemes must balance complexity against accuracy, and also take into account the fact that in some cases, there is no guaranteed solution. In those cases where we make a wrong bit decision, the bit will be in error, and must be corrected by some other means (in the case of orientation bits, the other corners may help, and in the case of the tag's data, the Reed-Solomon decoding stage will correct errors).

Before the data can be extracted from the data area, the pixel ranges for "black" and white dots needs to be ascertained. The minimum and maximum pixels encountered during the search for the tag were stored in `MinPixelEncountered` and `MaxPixelEncountered` respectively. The following pseudocode shows the method of obtaining good values for `WhiteMin` and `BlackMax` based on the min & max pixels encountered.

```

MinPixel = MinPixelEncountered
MaxPixel = MaxPixelEncountered
MidRange = (MinPixel + MaxPixel) / 2
WhiteMin = MaxPixel - PenConstantWhiteFactor
BlackMax = MinPixel + PenConstantBlackFactor

```

The scheme used to determine a macrodot's value if the pixel value is between `BlackMax` and `WhiteMin` is fairly simple, but gives good results. It uses the pixel values of the macrodot centers to the left and right of the dot in question, using their values to help determine a more likely value for the center dot:

- If the two macrodots to either side are on the white side of `MidRange` (an average macrodot value), then we can guess that if the center macrodot were white, it would likely be a "definite" white. The fact that it is in the not-sure region would indicate that the macrodot was black, and had been affected by the surrounding white macrodots to make the value less sure. The macrodot value is therefore assumed to be black, and hence the bit value is 1.
- If the two macrodots to either side are on the black side of `MidRange`, then we can guess that if the center macrodot were black, it would likely be a "definite" black. The fact that it is in the not-sure region would indicate that the macrodot was white, and had been affected by the surrounding black macrodots to make the value less sure. The macrodot value is therefore assumed to be white, and hence the bit value is 0.
- If one macrodot is on the black side of `MidRange`, and the other macrodot is on the white side of `MidRange`, we simply use the center macrodot value to decide. If the center macrodot is on the black side of `MidRange`, we choose black (bit value 1). Otherwise we choose white (bit value 0).

The logic is represented by this simple pseudocode:

```

if (pixel < WhiteMin)           // definitely black
    bit = 0x01
else
    if (pixel > BlackMax)       // definitely white
        bit = 0x00
    else

```

```

if ((prev > MidRange) && (next > MidRange))    //prob black
    bit = 0x01
else
if ((prev < MidRange) && (next < MidRange))    //prob white
    bit = 0x00
else
if (pixel < MidRange)
    bit = 0x01
else
    bit = 0x00

```

From this one can see that using surrounding pixel values can give a good indication of the value of the center macrodot's state. The scheme described here only uses the macrodots from the same row, but increased accuracy may be achieved using all the neighboring macrodots, at the cost of increased complexity.

We already have the distance between the macrodots in terms of a Δx and a Δy with respect to the current orientation of the image. Rather than rotate the image to transform Δx and Δy both to 1, we simply change the sense of the Δx and Δy values based on how many sets of 90 degrees it takes for the tag to be oriented so that both D values are positive. The four possible rotations by 90 degrees equate to 4 different combinations of signs ($\pm\Delta x$ and $\pm\Delta y$). It is therefore trivial to translate Δx and Δy into two deltas for movement in x and y: Δx_x and Δy_x for movement of 1 macrodot in x, and Δx_y and Δy_y for movement of 1 macrodot in y.

With the 12 orientation bits obtained, the orientation of the tag can be easily determined. Given 4 sets of 12 bits, the correct orientation is the one that minimizes the number of bits in error, as shown in Table 5.

Table 5. The four major orientations of the position tag

12-bit value	Orientation of tag
111-011-100-000	No rotation
011-100-000-111	Rotated clockwise 90 degrees
100-000-111-011	Rotated clockwise 180 degrees
000-111-011-100	Rotated clockwise 270 degrees

3.1.4.2 Extract the Tag's Mode Raw Data

With the orientation determined, the next step is to read the bit pattern representing the Mode bits from the tag.

We know the coordinate of the tag's target center as $[\text{Center}_x, \text{Center}_y]$. We also know the orientation of the tag and therefore have Δx_x and Δy_x for motion in x, and Δx_y and Δy_y for movement in y. These delta values enable us to step along the center of any macrodots. By sampling the tag at the 18 mode bit positions (see Figure 6). Consequently the location of the first few mode bit are:

- Bit0 location = $[\text{Center}_x - 5\Delta x_x - 5\Delta x_y, \text{Center}_y - 5\Delta y_x - 5\Delta y_y]$
- Bit1 location = $[\text{Bit0}_x + \Delta x_x, \text{Bit0}_y + \Delta y_x]$
- Bit2 location = $[\text{Bit1}_x + \Delta x_x, \text{Bit1}_y + \Delta y_x]$

Given the macrodot coordinate (fixed point) we use the sampling mechanism as described in Section 3.1.4.1 to extract the bit value.

The extraction of 28 bits will take approximately 168 cycles on a simple microprocessor. Assuming a clock speed of 25 MHz, 168 cycles per position sample represents less than 0.1% of the available processor bandwidth.

3.1.4.3 Decode the Mode Data

The 28-bit raw bitstream representing the tag's mode data has been read in. It must now be decoded into 8 bits of meaningful mode information.

The data is decoded using straightforward Reed-Solomon decoding. One of the advantages of using Galois Field GF(16) is that a total of only 64 bytes are required for lookup tables during the decoding process (16 bytes each of power, log, multiply and multiply inverse).

The decoding process of 28 bits to 8 bits will take approximately 1500 cycles on a simple microprocessor. Assuming a clock speed of 25 MHz, 1500 cycles per position sample is less than 1% of the available processor bandwidth.

3.1.4.4 Extract the Raw Tag Data

Before the tag data *proper* can be read in from the tag image, the tag's 8-bit mode data must be examined. The Version bit must be checked to ensure that this tag has a known structure (see Table 3). If the Version bit is 0, then the tag data structure is known and can be read in from the image.

The macrodots representing the tag data area are read into a 168-bit raw bitstream according to the locations described in Figure 7. We use the tag's target center [$Center_x$, $Center_y$] and the orientation delta values Δx_x and Δy_x , Δx_y and Δy_y to calculate the start position. The delta values also enable us to step along the center of the data area's macrodots.

The process of determining if a given macrodot is a 1 or a 0 is the same as that carried out for reading the tag's mode bits and orientation bits (see Section 3.1.4.1). A total of 168 bits is read in from the image in this way. The extraction of 168 bits will take approximately 1008 cycles on a simple microprocessor. Assuming a clock speed of 25 MHz, 1008 cycles per position sample represents 0.4% of the available processor bandwidth.

3.1.4.5 Decode the Tag Data

The 168-bit raw bitstream representing the tag data has been read in. It must now be decoded into the PageId, x and y coordinates etc. (see Table 2).

The 28×6 -bit codewords can be divided into message data codewords and check words. Assuming that the Version bit is 0, the 2-bit NumMsgWords field in the Mode area determines the number of message data codewords D as 15, 16, 17, or 18. The remaining K codewords are check codewords.

Each of the D 6-bit message codewords can now be examined to see if it is all 0s or all 1s. Any such codeword is illegal (due to the encoding process) and can be considered to have been erased. Determining that a codeword has been erased increases the error correction capabilities of the check bits.

The data can then be decoded using straightforward Reed-Solomon decoding. One of the advantages of using Galois Field GF(64) is that a total of only 256 bytes are required for

lookup tables during the decoding process (64 bytes each of power, log, multiply and multiply inverse).

The decoding process of 168 bits to 90 bits will take approximately 4000 cycles on a simple microprocessor. Assuming a clock speed of 25 MHz, 4000 cycles per position sample represents 1.6% of the available processor bandwidth.

3.1.5 Build Stroke Information

The information gained from the position tag is either ignored or added to the current stroke depending on whether the pen is deemed, based on nib pressure, to be up or down. If the pen is up but was down, then the current stroke is finished. If the pen was up but is now down, then a new stroke is started.

Each stroke contains the start time (in milliseconds) followed by the recovered page id and initial position. This is followed by a series of positions for the stroke. The positions are implicitly separated by 100ths of a second in time. While the first position is always absolute, subsequent positions are either delta-encoded or absolute as required. Escape codes allow compression for a number of unknown positions (the tag cannot be found), the encounter of different page ids (for example the initial page id is unknown and finally a page id is recovered from a tag during the stroke, or it may be that the pen crosses from one page to another during a stroke), and the encountering of tags with differing Pressure-SensitiveArea values, which enable or disable the inclusion of a pressure value with each position. Pressure values are typically included for the entire stroke if the stroke begins in a pressure-sensitive area.

The following is a definition for a stroke:

```
Stroke = time PageIdRec FirstPosRec {[PosRec | PressurePos | Escape]} endStroke
PageIdRec = [unknown | pageId]
Escape = [pageId | pressureMode | TimeSkip]
TimeSkip = [unknown | zeroDelta] shortTime
PressurePos = PosRec penPressure
PosRec = [PenAbsRec | PenDeltaRec]
FirstPosRec = [unknown | PenAbsRec]
PenAbsRec = tagCoordinate imageCoordinate macrodotDelta tilt
PenDeltaRec = penDelta imageCoordinate macrodotDelta tilt
```

where:

```
time = 32 bits unsigned (milliseconds)
shortTime = 16 bits unsigned (milliseconds)
pageId = 64 bits unsigned
tagCoordinate = 2 x 9 bits unsigned (representing x and y)
imageCoordinate = 2 x 10 bits unsigned fixed point 7:3 (representing x and y)
macrodotDelta = 2 x 8 bits signed fixed point S3:4 (representing Δx and Δy)
penDelta = 2 x 4 bits signed (representing -8..7 for x and y)
penPressure = 5 bits
tilt = 2 x 5 bits (representing 2 tilt axes)
pressureMode = 1 bit
```

The length of a sample position is 64 bits (18 + 20 + 16 + 10) when the pen coordinate is fully described, and 52 bits (18 + 8 + 16 + 10) when the pen coordinate is a delta amount. Adding flag bits to specify that there is no escape mode (1 bit) and to select between absolute and relative position (1 bit), we specify 66 and 54 bits respectively.

The initial position requires 162 bits to define (32 + 65 + 65).

For a stroke over a single page that lasts 4 seconds, and assuming no read errors, there is an initial pen down followed by 399 deltas. The total stroke length is therefore 21,708 bits ($162 + (399 \times 54)$), which is 2714 bytes, or 6.8 bytes per sample.

If pressure is also captured with each position (depending on the value of the encountered PressureSensitiveArea bits within each tag), a further 5 bits are required for each position. The Escape of pressureMode allows the toggling of this state within a stroke.

The stroke definition copes well with unknown pen locations. If the initial position (pageId, coordinates etc.) is unknown due to the inability to locate a tag or recover tag data, the initial position for the stroke is marked as unknown. This initial position may be followed by an Escape of TimeSkip, signifying the amount of time where unknown positions were captured. Finally, when a tag is captured and decoded successfully, the Escape for pageId is used. This pageId is valid until the next Escape pageId is included in the stroke (inserted when a tag is encountered with a pageId different to the current *known* pageId). If a tag cannot be found or decoded successfully mid-stream within a stroke, the Escape for TimeSkip is used, which represents the elapsed time where only unknown positions were sampled. An unknown pen position does not change the pageId. If tag positions are reacquired, the page id may be changed, depending on the tag. A Timeskip entry for a set of unknown positions uses 1 bit for the escape code, 2 bits to define the type of escape code, and 16 bits to define the duration of samples skipped.

The other use for the Escape of TimeSkip is when there have been one or more deltas of 0. The Escape of zeroDelta combines with a 16-bit time amount to specify how long the pen has remained in the same position without moving.

The pen contains 0.5MB of RAM to buffer stroke information both during stroke construction and if the pen is out of range of the printer. A 0.5MB buffer allows for approximately 192 strokes at 4 seconds per stroke, which represents more than 12 minutes of stroke capture. The exact duration will depend on the success of tag decoding, and whether or not continuous pressure is being captured.

If the stroke buffer becomes full, a *Buffer-Full* LED is illuminated. The LED stays illuminated until the pen comes into range of a printer, whereupon the buffered strokes are transmitted to the printer and the buffer freed.

3.1.6 Activate Active Area Light

The position tag data area contains a ActiveArea bit (see Table 2). The pen's active area LED must be enabled directly in response to the value of this bit. There is no additional interpretation of the ActiveArea bit by the pen.

3.2 DATA ENCRYPTION

The data stream must be encrypted before being transmitted to the printer, to prevent eavesdroppers from capturing pen input, and particularly signature input.

Consequently, the final stage before transmission is the encryption of the data stream. Symmetric encryption is used since it is tractable in the time available. Possible symmetric

algorithms for use are listed in Table 6. In all cases, a minimum key size of 14 bytes is assumed.

Table 6. Symmetric encryption algorithms [5]

Algorithm	Type	Clocks per byte processed (Pentium)	Block size (bytes)	Key size (bytes)	Patents
Blowfish	Block	18	8	4-56	No
CAST	Block	20	8	16	Yes, but free
IDEA	Block	50	8	16	Yes
Triple-DES	Block	108	8	14 ^a	No

a. Although Triple-DES uses a 168 bit key, the effective key strength is only 112 bits

Of the algorithms listed in Table 6, Triple-DES is the most conservative choice, although it is the slowest. At 108 clocks per byte processed, the encoding of a single 8-byte block consumes 864 cycles (8×108).

The pen must send 100 encrypted samples each second. Assuming that each sample requires a complete Triple-DES block (8 bytes), the time taken to encrypt 1 second of data is 86,400 cycles.

For a clock speed of 25 MHz, 86,400 cycles for the 100 samples equates to less than 0.5% of the processor bandwidth.

The encryption process therefore has a throughput of 800 bytes per second.

REFERENCES

4 Silverbrook References

5 Other References

- [1] Fishkin, K., "Filling a Region in a Frame Buffer", *Graphics Gems*, Academic Press 1990, pp.278-284
- [2] Heckbert, P.S., "A Seed Fill Algorithm", *Graphics Gems*, Academic Press 1990, pp.275-277
- [3] Lypkens, H., "Reed-Solomon Error Correction", *Dr. Dobbs's Journal* Vol.22, No.1, January 1997
- [4] Rorabaugh, C, *Error Coding Cookbook*, McGraw-Hill 1996
- [5] Schneier, S., "The Crypto Bomb is Ticking", *Byte Magazine*, May 1998, pp.97-102
- [6] Turkowski, K., "Filters for Common Resampling Tasks", *Graphics Gems*, Academic Press 1990, pp.147-165
- [7] Welch Allyn Inc., *Aztec Barcode Symbology Specification Rev 3.0*, 1995
<http://dcd.welchallyn.com/techover/dcdwhite.htm>
- [8] Wicker, S., and Bhargava, V., *Reed-Solomon Codes and their Applications*, IEEE Press 1994

Netpage Applications

draft version 0.2, 20 October 1999



Silverbrook Research Pty Ltd
393 Darling Street, Balmain
NSW 2041 Australia
Phone: +61 2 9818 6633
Fax: +61 2 9818 6711
Email: info@silverbrook.com.au

Confidential

Contents

APPLICATION ARCHITECTURES	1
1 Netpage Applications	2
1.1 Object Model Description	2
1.2 User Interface Description	3
1.3 Netpage User Interface Design Guidelines	4
2 Directory Navigation	6
3 User Registration	9
3.1 User Object Model	9
3.2 User Registration User Interface	10
4 Electronic Mail	22
4.1 E-mail Object Model	22
4.2 E-mail User Interface	25
5 Purchasing	53
5.1 Purchasing Object Model	53
5.2 Purchasing User Interface	58
6 Folder Hierarchy	65
6.1 Folder Hierarchy	65
7 Navigation Aids	66
7.1 Favorites List	66
7.2 History List	67
8 Drawing and Painting	68
8.1 Lasso	68
9 Viewing 3D Objects	69
AUTHORING	71
10 Authoring Transition Strategies	72
10.1 Overview	72
10.2 Re-purposing Editorial Content	72
10.3 Re-purposing Web Applications	73
10.4 Re-purposing Web Content	73
REFERENCES	74
11 Silverbrook References	75
12 Other References	76

APPLICATION ARCHITECTURES

1 Netpage Applications

The Netpage System can be used to deliver many of the same interactive computer applications normally delivered via networked computer terminals such as Web terminals.

In the Web paradigm, a series of Web “pages” displayed on the terminal screen provide the user interface to the application. In the Netpage paradigm, a series of printed Netpages provide the user interface to the application.

The following sections define a set of core Netpage applications, including user and pen registration, news subscription, electronic mail (e-mail), and various examples of electronic commerce (e-commerce). In each case the application is described from the point of view of the user, in terms of the Netpage user interface seen by the user and the object model implicitly manipulated by the user through the user interface.

1.1 OBJECT MODEL DESCRIPTION

Each object model is described using a Unified Modeling Language (UML) class diagram [3,9]. A class diagram consists of a set of object classes connected by relationships. Two kinds of relationships are of interest here: associations and generalizations. An association represents some kind of relationship between objects, i.e. between instances of classes. A generalization relates actual classes, and can be understood in the following way: if a class is thought of as the set of all objects of that class, and class A is a generalization of class B, then B is simply a subset of A. The UML does not directly support second-order modelling - i.e. classes of classes.

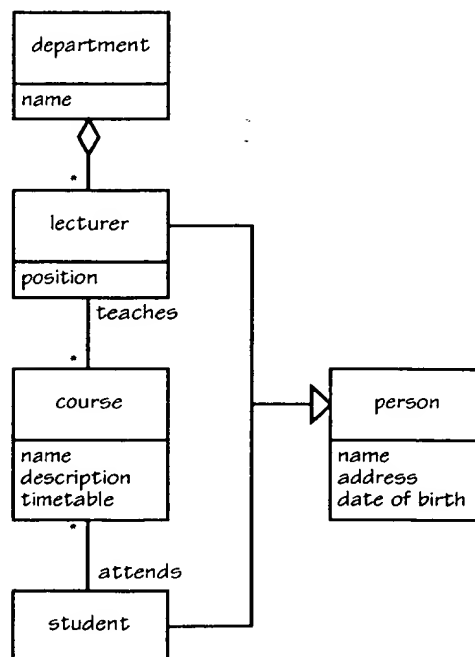


Figure 1. Illustrative UML class diagram

Each class is drawn as a rectangle labelled with the name of the class. It contains a list of the attributes of the class, separated from the name by a horizontal line, and a list of the

operations of the class, separated from the attribute list by a horizontal line. In the class diagrams which follow, however, operations are never modelled.

An association is drawn as a line joining two classes, optionally labelled at either end with the multiplicity of the association. The default multiplicity is one. An asterisk (*) indicates a multiplicity of "many", i.e. zero or more. Each association is optionally labelled with its name, and is also optionally labelled at either end with the role of the corresponding class. An open diamond indicates an aggregation association ("is-part-of"), and is drawn at the aggregator end of the association line.

A generalization relationship ("is-a") is drawn as a solid line joining two classes, with an arrow (in the form of an open triangle) at the generalization end.

When a class diagram is broken up into multiple diagrams, any class which is duplicated is shown with a dashed outline in all but the main diagram which defines it. It is shown with attributes only where it is defined.

Figure 1 illustrates the subset of the UML class diagram notation used here.

1.2 USER INTERFACE DESCRIPTION

Each application user interface flow is illustrated as a collection of documents linked by command arrows. A command arrow indicates that the target document is printed as a result of the user pressing the corresponding command button on the source page. Some command arrows are labelled with multiple commands separated by slashes ('/'), indicating that any one of the specified commands causes the target document to be printed. Although multiple commands may label the same command arrow, they typically have different side-effects.

In application terms, it is important to distinguish between Netpage documents and Netpage forms. Documents contain printed information, as well as command buttons which can be pressed by the user to request further information or some other action. Forms, in addition to behaving like normal documents, also contain input fields which can be filled in by the user. They provide the system with a data input mechanism.

It is also useful to distinguish between documents which contain generic information and documents which contain information specific to a particular interaction between the user and an application.

Each document type has a distinct iconic representation as shown in Figure 2.




Figure 2. Document icons used in user interface flows

Generic documents may be pre-printed publications such as magazines sold at news stands or advertising posters encountered in public places. Forms may also be pre-printed, including, for example, subscription forms encountered in pre-printed publications. They

may, of course, also be generated on-the-fly by a Netpage Printer in response to user requests.

User-specific documents and forms are normally generated on the fly by a Netpage Printer in response to user requests.




Figure 3. Page layout elements

The Netpages which participate in a user interface flow are further described as abstract page layouts. A page layout may contain fixed information, variable information, input fields, command buttons, and lists of draggable objects and attributes. Each type of layout element has a unique style to differentiate it from the others, as illustrated in Figure 3.

When a user interface flow is broken up into multiple diagrams, any document which is duplicated is shown with dashed outlines in all but the main diagram which defines it.

1.3 NETPAGE USER INTERFACE DESIGN GUIDELINES

Like a Web application, a Netpage application is forms-oriented and transaction-driven. The application is largely decoupled from the user interface in the sense that it is only notified when the user submits a form, not while the user is filling in the form. When the time it takes to deliver a form is short, a transaction-driven system can approximate the behavior of a truly interactive system. Such is often the case with the Web. With the Netpage System, however, the time it takes to deliver a form is longer, but more importantly, a form is not delivered *in situ*, but in the shape of another physical piece of paper. The new page must be physically collected from the printer, which may not be proximate to the user, and even if the actual printing cost of the page is small, the user may perceive the user interface as being wasteful.

It therefore behooves a good Netpage application user interface to be more sparing in the printing of intermediate pages than an equivalent screen-based application, and to ensure that the content of each printed page is valuable to the user whenever a page is printed.

1.3.1 Input Feedback

As with any computer application, it is important to provide the user with immediate feedback on all input, so that the evolving state of the transaction is visible and not hidden.

When a user fills in a Netpage form, immediate feedback is provided because the user uses the marking nib of the Netpage Pen. This applies to input in textboxes, checkboxes, signature fields and the like. When the user presses a command button, there is no immediate feedback on the page since the user uses the non-marking nib of the pen, but feedback is

nonetheless provided in the shape of the next form in the user interface series, or a transaction receipt, or some other requested document.

The user also uses the marking nib of the pen to input commands which don't require immediate interpretation, such as text editing markups.

The following are examples of input typically made with the marking nib:

- handwritten text
- drawing
- checkbox mark
- signature
- text editing command (strikeout etc.)
- drag-and-drop (from list of objects, etc.)
- drag-and-lasso (from palette of colors, line styles, etc.)

The following are examples of input typically made with the non-marking nib:

- selection by circumscription
- button press

1.3.2 A Minimum Unit of User Interface Output

A Netpage Printer prints at a constant page rate, independent of whether pages are printed on both sides or not, and independent of the pages' actual content.

A single page printed on both sides at high resolution in full color is therefore a useful minimum unit of user interface output. This minimum unit of output has an information-bearing capacity equivalent to ten SVGA computer screens. So although a Netpage application is designed to produce less frequent output than the equivalent screen-based application, it has ten times the "display" area of the screen-based application when it does produce output.

Furthermore, previous pages printed as part of the user interface flow continue to act as part of the user interface, and multiple applications can be active simultaneously without competing for screen space and forcing the user to switch between them on the screen.

2 Directory Navigation

A directory is an organized list of named objects. It may be presented in many ways: sorted alphabetically, arranged into named groups or topics, etc. The default presentation is specific to each kind of directory. A directory may also be searched.

For the purposes of navigation, an n -ary index tree is built above the directory, with each node containing enough index entries to fill two printed pages (i.e. a double-sided sheet). If two pages hold k entries, and the entire directory contains m entries; then the number of levels in the index is $\log_k m$. For example, if two pages hold 200 entries, and the entire directory contains 1,000,000,000 entries, then the number of levels in the index is 4.

Each index entry specifies the starting and ending names on the two pages it is linked to, separated by "to", e.g. "Aardvark to Axolotl", "Ayatollah to Bernoulli", etc. Each index entry acts as a hyperlink to the corresponding lower-level index node. The layout of the generic directory index page is shown in Figure 7.

If the entire directory fits on two pages, then it is presented directly without an intervening index tree.

The indexed directory class diagram is shown in Figure 20.

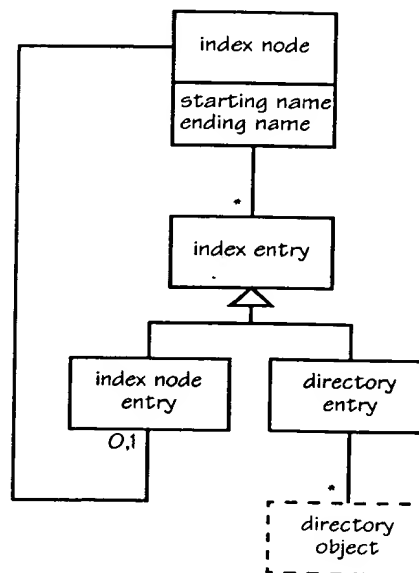


Figure 4. Indexed directory class diagram

The presentation of the index entries of a terminal (bottom-level) index node, i.e. a node which refers to actual named objects, is specific to each kind of directory.

Each page contains a <top> button which prints the root node of the index, an <up> button which prints the page's parent node in the index, and <first>, <previous>, <next> and <last> buttons which print the first, previous, next and last nodes, respectively, of the page's index level. Particular navigation buttons don't appear if they have no meaning for the current page, e.g. the page for the first index node in an index level doesn't have <first> and <previous> buttons.

Each page also contains a <search> button which generates a search form which allows the directory to be searched. The contents of the search form are specific to each kind of directory.

The directory navigation user interface flow is illustrated in Figure 5.

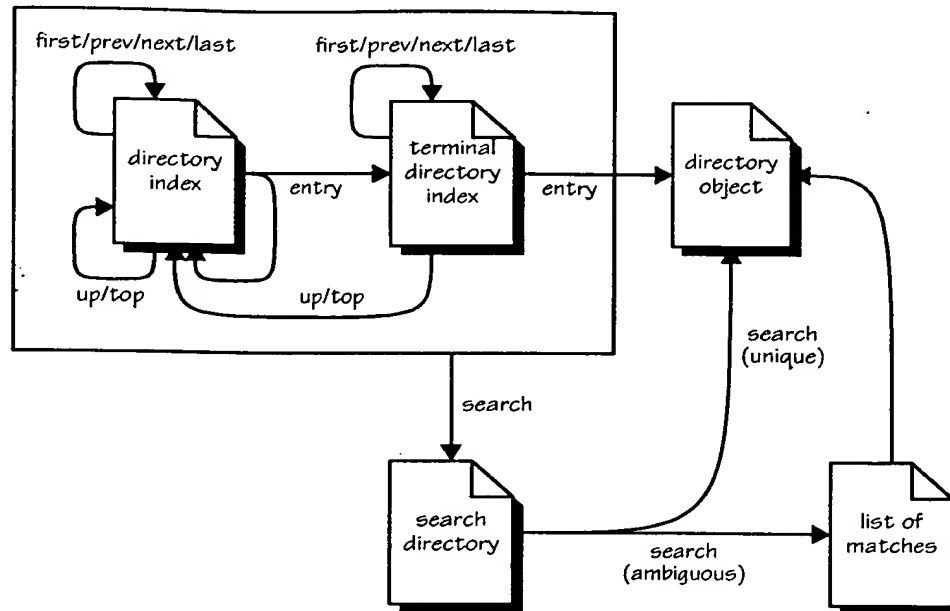


Figure 5. Directory navigation user interface flow

The meaning of the standard navigation buttons on the directory index page is described in Figure 6.




Figure 6. Standard directory index navigation button

[illegible]

Figure 7. Generic directory index page

3 User Registration

Much of the Netpage System is predicated on the system's secure knowledge of user identity. This allows the system to:

- deliver personalized services to individual users independent of location
- support authenticated transactions
- support secure payments
- guarantee a high level of privacy

3.1 USER OBJECT MODEL

The user object model revolves around a uniquely identified user, as shown in the class diagram in Figure 8.

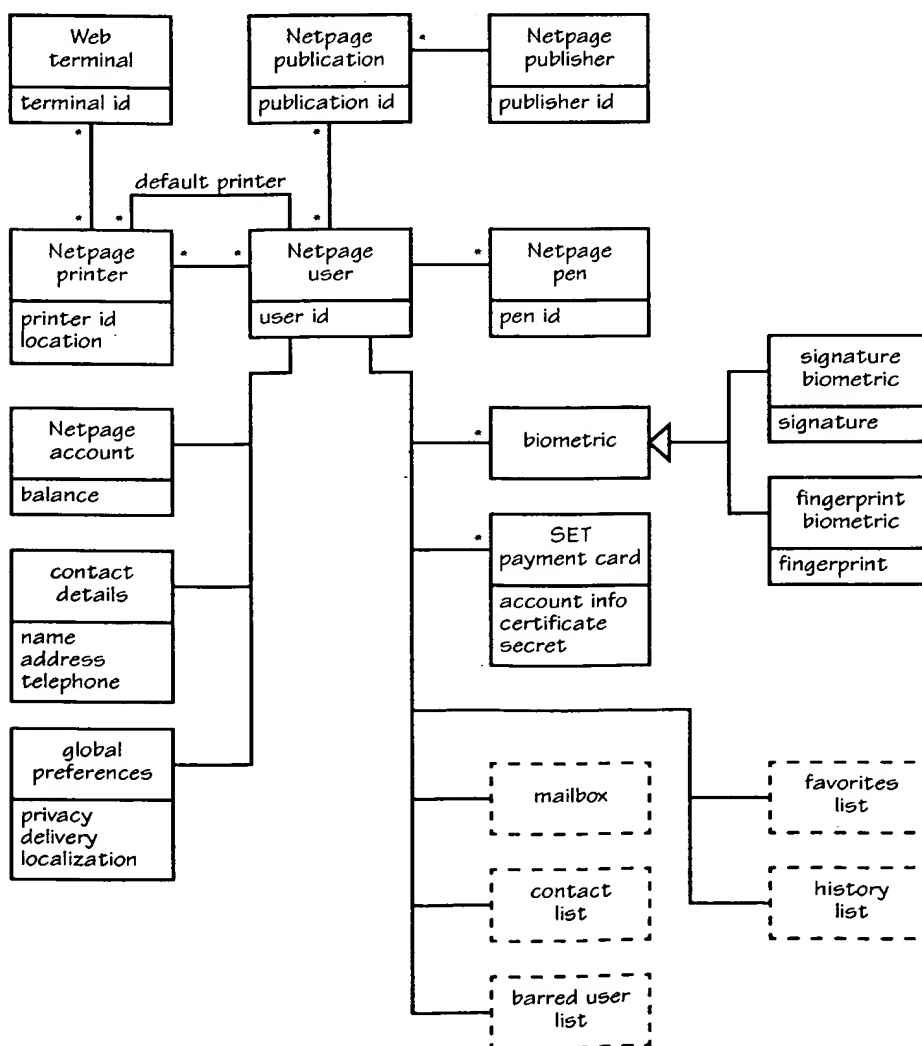


Figure 8. User class diagram

A user is authorized to use any number of printers, and printer may have any number of users. One of the user's printers is designated as the default printer, and subscriptions are automatically delivered to the default printer. Pages requested ad hoc, on the other hand, are automatically delivered to the printer through which the user is interacting. Thus a publisher never records the id of a printer, but instead resolves the id when it is required.

When a user subscribes to a publication, the publisher is authorized to print to the user's default printer. This authorization can be revoked at any time by the user.

A Web terminal may be authorized to print to a Netpage Printer via the Netpage Network. This allows it to conveniently print Web and Netpage documents encountered while navigating the Web.

A user may have any number of pens, but each pen is linked to a particular user. Each of a user's pens may be used interchangeably.

Each user also has a Netpage account which is used to accumulate micro-transactions; contact details, including name, address, and telephone numbers; global preferences, including privacy, delivery, and localization settings; any number of biometric records, containing the user's encoded signature, fingerprint etc.; and any number of payment card accounts, with which e-commerce payments can be made.

The user's mailbox and contact list are described in Section 4.

3.2 USER REGISTRATION USER INTERFACE

3.2.1 Help Page

The help page contains the following buttons which generate registration forms:

- add user
- add pen to user
- add pen to local user
- authorize user on printer
- authorize global user on printer
- authorize terminal on printer

3.2.2 User Registration

A user must be registered with the Netpage Network, own at least one registered Netpage Pen, and be authorized on at least one Netpage Printer, before being able to meaningfully interact with the Netpage Network.

The <add user> button on the help page generates a **user registration** form (Figure 10) which captures the user's contact details and privacy preferences.

The pen used to submit the user registration form, if previously unregistered, is automatically linked to the new user. The user is automatically authorized to use the printer through which the registration takes place.

A subset of the users authorized to use a printer are designated as administrators, and only they may authorize the registration of additional users and pens on the printer.

By default, every new user is an administrator on a particular printer until a printer-specific privileges setting is enabled (via the help page). Thus a new user can be registered without having to be authorized by an administrator's pen until the privileges setting is enabled. This setting may never need to be enabled in a domestic environment, but may commonly be enabled in a corporate environment. User-specific privileges settings only appear on the registration form if the printer-specific privileges setting is enabled.

When privileges are enabled, the user registration form must be authorized by an administrator's pen. This is done by initialling the registration form in the appropriate place.

The user registration user interface flow is illustrated in Figure 9.




Figure 9. User registration user interface flow

Add Netpage User

page 1 of 1

Title

title

First Name

first name

Initials

initials

Nickname

nickname

Family Name

family name

Street Address

street address 1

street address 2

City

city

State

state

ZIP Code

ZIP Code

Country

country

Work Phone

work phone number

Mobile Phone

mobile phone number

Home Phone

home phone number

Privacy Options

☐ Don't publish name in global user list

☐ Don't publish address

☐ Don't publish phone numbers

☐ Only accept e-mail from known contacts

Privileges

☐ Administrator

Initialed by Administrator

initials

submit

Figure 10. User registration page

3.2.3 Pen Registration

Additional Netpage Pens may be linked to existing users.

The <add pen to user> button on the help page generates an **pen registration form** (Figure 12). The form contains name and nickname fields and a <submit> button. If the form is requested using the new pen, i.e. the pen to be registered, then the form is blank and must be filled in with sufficient name details to identity the user. If the form is requested using an existing pen, then the form is automatically filled in with the details of the existing pen's owner.

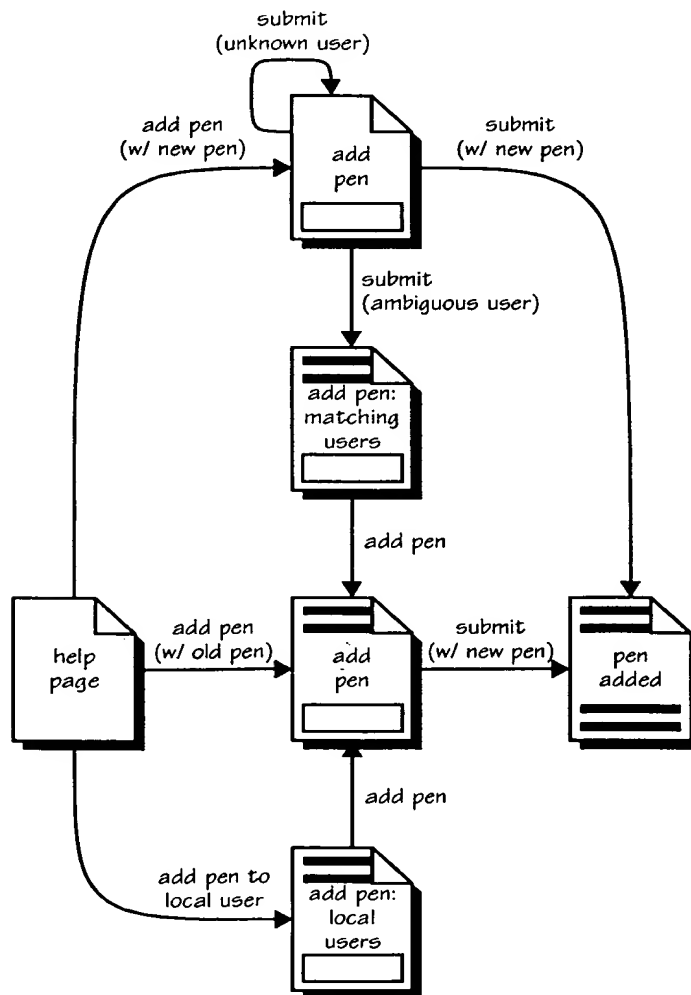


Figure 11. Pen registration user interface flow

When the form is submitted, the name information is matched against the list of local users, i.e. users authorized to use the printer. Partial name information can be entered, such as just a first name or a nickname. If the name information is ambiguous, the system generates a list of matching users (Figure 13) with an <add pen> button next to each.

The user must submit the pen registration form using the new pen. This links the new pen to the user. The form cannot be submitted using a pen which is already registered.

The <add pen to local user> button on the help page generates a list of all users authorized to use the printer (Figure 14), with an <add pen> button next to each. The local user list is only available to administrators.

When privileges are enabled, the pen registration form must be authorized by an administrator's pen. This is done by initialling the registration form in the appropriate place.

The pen registration user interface flow is illustrated in Figure 11.

The form is titled "Add Pen to Netpage User" in a cursive font at the top left, and "page 1 of 1" at the top right. It contains several input fields: "First Name" with a box labeled "first name", "Initials" with a box labeled "initials", "Nickname" with a box labeled "nickname", and "Family Name" with a box labeled "family name". At the bottom right, there is a label "Initialed by Administrator" next to a box labeled "initials", and a rounded "submit" button below it.

Figure 12. Pen registration page

*Add Pen to Netpage User:
Matching User List*

[page # of #]

First Name	first name	Initials	initials
Nickname	nickname		
Family Name	family name		

add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name
add pen	user name

Figure 13. Pen registration - matching users page

[illegible]

Figure 14. Pen registration - local users page

3.2.4 Authorize User on Printer

A user may be authorized to use additional Netpage Printers. Once authorized to use a particular printer, the user may interact with the printer using any pen linked to the user.

The <authorize user on printer> button on the help page generates a **user authorization** form (Figure 16). The registration form contains name and nickname fields and a <submit> button. If the form is requested using a pen owned by a local user, then the form is blank and must be filled in with sufficient name details to identity the new user. If the form is requested using the new user's pen, then the form is automatically filled in with the details of the new user.

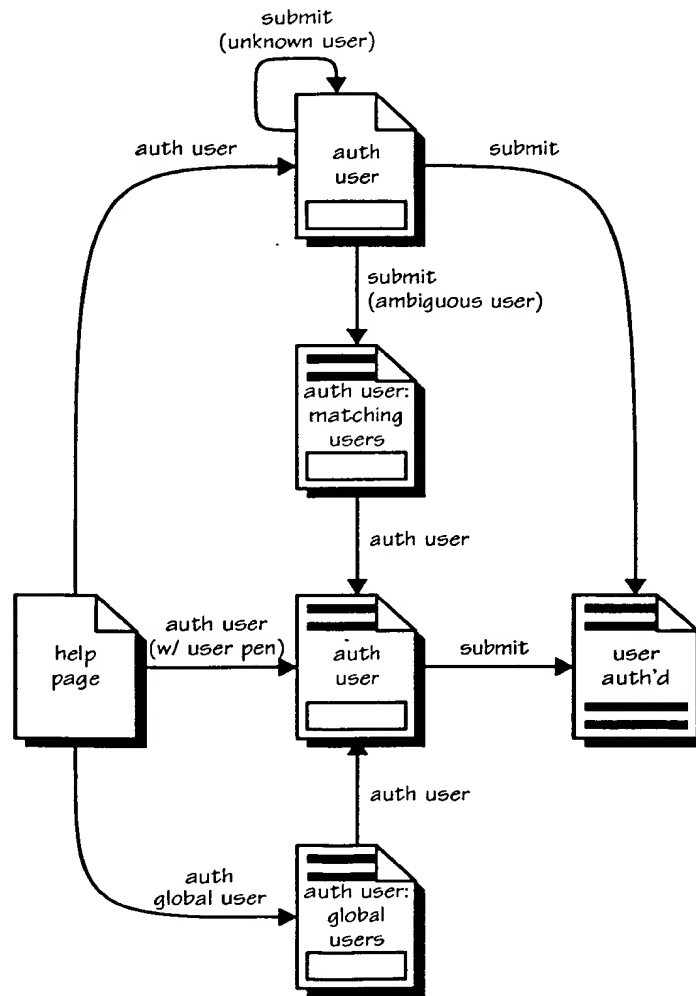


Figure 15. User authorization user interface flow

When the form is submitted, the name information is matched against the global user list. Partial name information can be entered, such as just a first name or a nickname. If the name information is ambiguous, the system generates a list of matching users (Figure 17) with an <authorize user> button next to each.

The <authorize global user> button on the help page generates a list of all users on the network (Figure 18), with an <authorize user> button next to each. The global user list is available to all users, but only contains users who have elected to appear.

When privileges are enabled, the user authorization form must be authorized by an administrator's pen. This is done by initialling the authorization form in the appropriate place.

The user authorization user interface flow is illustrated in Figure 15.

*Authorize Netpage User on Printer**page 1 of 1*

First Name

first name

Initials

initials

Nickname

nickname

Family Name

family name

*Initialed by
Administrator*

initials

submit

Figure 16. User authorization page

Authorize Netpage User On Printer - Matching User List page # of #

First Name	<div>first name</div>	Initials	<div>initials</div>
Nickname	<div>nickname</div>		
Family Name	<div>family name</div>		

<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>
<div>auth user</div>	<div>user name</div>

Figure 17. User authorization - matching users page

[illegible]

Figure 18. User authorization - global users page

3.2.5 Authorize Terminal on Printer

The <authorize terminal on printer> button on the help page produces a terminal authorization token. The token is printed with information which allows the user of the Web terminal to register the Web terminal with the Netpage Network, allowing it to print to the Netpage Printer identified by the token.

An administrator may revoke a terminal's authorization at any time. The <authorize terminal on printer> button is also only available to administrators.




Figure 19. Terminal authorization user interface flow

4 Electronic Mail

Netpage electronic mail (e-mail) provides a messaging service between Netpage users. It also supports message exchange with Internet e-mail users, and by extension, users of other e-mail systems interconnected with the Internet, such as corporate e-mail systems.

Each Netpage user has a unique id within the Netpage System, and is also known to the system by their name, which may or may not be unique. Netpage e-mail is usually addressed by selecting a name from a list, for example from the global list of Netpage users or from a particular user's list of contacts. The user's nickname or alias helps disambiguate similar names. A Netpage user usually doesn't have to know or specify another Netpage user's unique id.

When a Netpage user wants to send e-mail to an Internet user, they must specify an Internet e-mail address. Similarly, when an Internet user wants to send e-mail to a Netpage user, they must specify the Netpage e-mail address. A Netpage e-mail address takes the form `<userid>@netpage.net`.

E-mail is a core service of the Netpage System.

4.1 E-MAIL OBJECT MODEL

4.1.1 E-mail User

The e-mail object model revolves around an e-mail user. An e-mail user is either a Netpage user or an Internet user.

The e-mail user class diagram is shown in Figure 20.




Figure 20. E-mail user class diagram

4.1.2 E-mail

The e-mail itself consists of a number of pages. Each page corresponds to a Netpage. The page structure is logical because the e-mail is delivered in the same form that it is composed, and it is composed by the sender page by page.

Each page contains digital ink, together with any number of attachments. An attachment is printed at its point of insertion, and may overflow onto the following page(s).

The sender of the e-mail is an e-mail user, as are any number of recipients and copy recipients (CCs). Blind copy recipients (BCCs) are not modelled here for clarity, but can obviously be trivially accommodated.

The e-mail has an explicit subject which is always converted from digital ink to text to support interoperability with e-mail systems which don't directly support digital ink, such as e-mail systems on the Internet, and to streamline the presentation of mailbox contents within the Netpage System itself.

The e-mail has a high priority flag which allows the both sender and the recipients to control the timeliness of its delivery.

The e-mail class diagram is shown in Figure 21.




Figure 21. E-mail class diagram

4.1.3 Mailbox

Incoming e-mails may be accumulated in the user's mailbox; printed on the user's default printer, incorporated into the user's daily newspaper, or any combination of these. The user can select, for example, to print high-priority e-mails immediately on receipt but to hold low-priority e-mails.




Figure 22. Mailbox class diagram

The user may create any number of named folders within the mailbox, and either manually copy and move received e-mails between folders, or associate a folder with one or more e-mail contacts (including e-mail groups) to have e-mail received from those contacts automatically placed in the folder. E-mail received from contacts not associated with a folder is placed in a predefined "inbox" folder. All e-mail sent by the user is placed in a predefined "sent e-mail" folder for future reference.

Each folder has three delivery options: "print all e-mail", "print high-priority e-mail", and "delete e-mail once printed". If a print option is selected, e-mail messages with the corresponding priority are printed on the user's default printer immediately on receipt. If the delete option is selected, e-mail messages are deleted from the e-mail folder once printed. Otherwise e-mail messages are held indefinitely in the e-mail folder, i.e. until manually deleted by the user.

The mailbox class diagram is shown in Figure 22.

4.1.4 Contact List

Each user has a list of contacts. The contact list provides a more convenient basis for selecting e-mail recipients than the global list of users, particularly since most user's contact lists will fit on a single double-sided printed page. The contact list also provides a basis for ignoring unsolicited e-mail. A privacy option allows a user to ignore e-mail not sent by a member of the contact list.

A user may create any number of contact groups within the contact list, and treat a group as a single contact for the purposes of addressing outgoing e-mail and controlling the delivery of incoming e-mail.

Groups may themselves contain groups, and both individual contacts and groups may be members of multiple groups as well as the top-level contact list itself.

The contact list class diagram is shown in Figure 23.

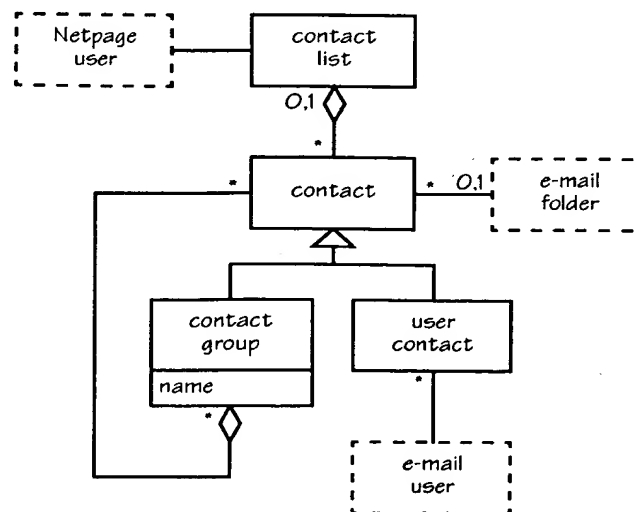


Figure 23. Contact list class diagram

4.1.5 Barred User List

Rather than accepting e-mail only from known contacts, a user may choose to bar individual users. The barred user list records individuals from which the user refuses to accept e-mail.

The barred user list class diagram is shown in Figure 24.




Figure 24. Barred user list class diagram

4.2 E-MAIL USER INTERFACE

4.2.1 Send E-mail

E-mail is sent using the **outgoing e-mail** form (see Figures 27 and 28). The e-mail form is always printed with the sender's name already specified, since the sender's identity is known from the pen used to request the e-mail form. If the e-mail form is requested by pressing a user-specific <e-mail> button on some page, then the e-mail form is printed with the recipient's name already specified. On the other hand, if the e-mail form is requested from the help page, then the recipient is unknown and is left blank on the form. The user must specify at least one recipient before the e-mail can be sent.

The user can also pre-address an e-mail form by pressing the <e-mail to contacts> button on the help page. This elicits an **add recipients** form (Figure 29), which lists the user's contacts. Each contact has a <To> and a <CC> checkbox as well as an <e-mail> button. The e-mail button elicits an e-mail form addressed to that user. The <e-mail selected> button at the bottom of the page elicits an e-mail addressed to all the users whose <To> or <CC> checkboxes have been checked.

The e-mail form initially consists of a double-sided page. The front (Figure 27) contains fields for the names of recipients and copy recipients, for the subject, and for the body of the e-mail. The back (Figure 28) contains a field for continuing the body of the e-mail. Any recipient names (or addresses) written by hand are converted from digital ink to text for lookup purposes. The subject is also converted, for presentation purposes, as discussed earlier. The body is retained as digital ink, to allow handwritten text and diagrams etc. to be delivered in a uniform and expressive manner.

The <add page> button at the bottom of every page of the e-mail form adds another double-sided page to the e-mail form. Only the additional page is printed, not the entire form,

but the additional page is logically linked to the original e-mail. The <discard page> button at the bottom of every page discards the corresponding page from the e-mail (but not both sides of the double-sided page). The entire e-mail is reprinted with the discarded page removed and subsequent pages, if any, moved to close the gap. A blank page added to make the page count even, if necessary.

Each printed e-mail form corresponds to a separate e-mail instance, the uniqueness of which is indicated by the sender's name, date and time printed at the top of every page. Once an e-mail has been sent, it can't be sent again. A copy, however, can be easily made, edited, and sent. The standard <print> button on every page of the e-mail form prints another copy of the e-mail form, corresponding to a new e-mail instance. Both the original and the copy can be edited further and sent independently.

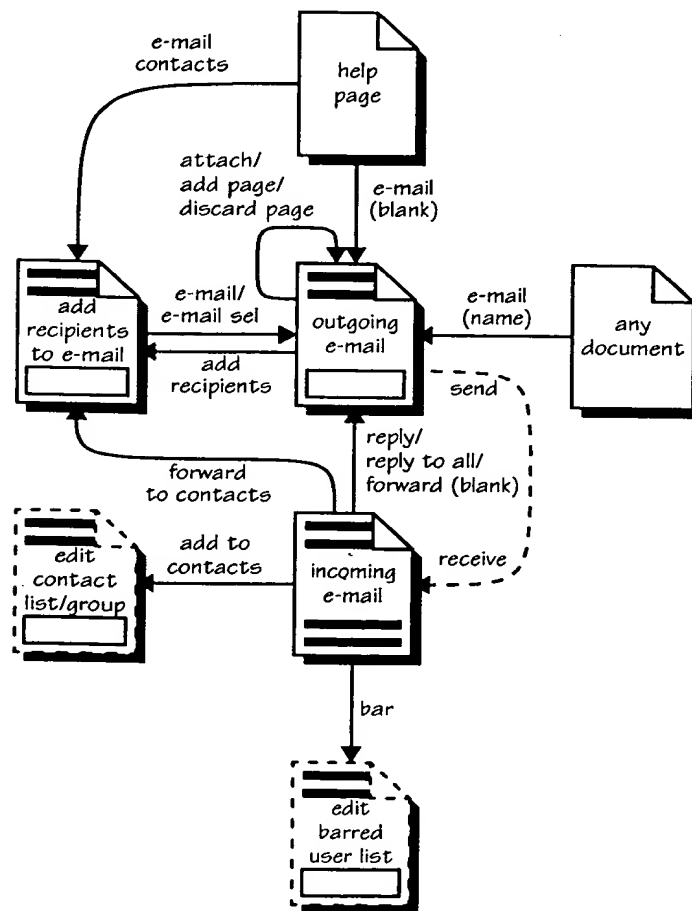


Figure 25. E-mail user interface flow

Text in the <recipients>, <copy recipients> and <subject> can be struck out with the pen to remove recipients, subject text, etc. The <print> button produces a copy of the e-mail form with the struck-out text removed.

The <attach> button at the bottom of every page attaches the current selection at the current end of the e-mail body. The entire e-mail is reprinted with the attachment included.

Additional pages are automatically added to the e-mail to accommodate the attachment. The attachment can consist of anything which can be selected on any Netpage.

The <add recipients> button adjacent to the <recipients> and <copy recipients> fields at the top of the first page of the e-mail form elicits an add recipients form (Figure 29), with the <subject> field reflecting the subject of the e-mail form. The <e-mail> and <e-mail selected> buttons on the add recipients form elicits a copy of the e-mail form with the additional recipients and copy recipients added.

The <send> button at the bottom of every page sends the entire e-mail. If the name of any recipient is unknown to the system, then the e-mail form is reprinted with the offending name printed in red, together with an error message indicating the problem.

If the name of any recipient is ambiguous, then a list of matching users is printed, each with a checkbox allowing it to be selected. The <e-mail> button at the bottom of the form reprints the e-mail form with the recipient names suitably updated.




Figure 26. E-mail recipient exception user interface flow

The overall e-mail user interface flow is illustrated in Figure 25. The e-mail recipient exception user interface flow is illustrated in Figure 26.

E-mail from sender at time & date *page # of #*

To High Priority

CC

Subject

☐ line color palette ☐ line width palette

e-mail body

Figure 27. Outgoing e-mail first page

E-mail from sender to recipient at time & date

page # of #

line color palette

line width palette

e-mail body

attach

add page

discard page

send

Figure 28. Outgoing e-mail second or subsequent page

[illegible]

Figure 29. Add recipients page

4.2.2 Receive E-mail

E-mail is received in the form of an **incoming e-mail** document (Figures 30 and 31), which has the same page structure and content as the corresponding outgoing e-mail form.

The <reply> button at the bottom of every page of the e-mail document produces an outgoing e-mail form (Figures 27 and 28) addressed to the sender of the incoming e-mail, and with the subject reflecting the subject of the incoming e-mail, but with "In reply to:" added as a prefix.

The <reply to all> button produces an outgoing e-mail form addressed to the sender of the incoming e-mail, as well as to all of its recipients and copy recipients.

The <forward> button produces an outgoing e-mail form with no recipient, and with the subject reflecting the subject of the incoming e-mail, but with "Forwarded:" added as a prefix. The body of the incoming e-mail is also copied to the body of the outgoing e-mail.

The <forward to contacts> creates, but doesn't print, an outgoing e-mail in the same way as the <forward> button, and then implicitly invokes the <add recipients> command on it. This elicits an add recipients form in the usual way (Figure 29), with the <subject> field reflecting the subject of the forwarded e-mail.

The <add to contacts> button adds the sender to the recipient's list of contacts and produces an **edit contacts** form reflecting the updated contact list. Contact list editing is described in more detail in Section 4.2.3 on page 34.

The <bar> button adds the sender to the recipient's list of barred users and produces an **edit barred users** form reflecting the updated barred user list. Barred user list editing is described in more detail in Section 4.2.5 on page 43.

E-mail for recipient at time & date *page # of #*

From sender

To recipient(s)

CC copy recipient(s)

Subject subject

e-mail body

Figure 30. Incoming e-mail first page

The diagram illustrates the layout of an incoming email interface. It features a large rectangular frame with a dashed border. At the top left, a dashed box contains the text "E-mail for recipient from sender at time & date". At the top right, another dashed box contains the text "page # of #". Below these, the main area is labeled "e-mail body". At the bottom of the frame, there are four rounded rectangular buttons labeled "reply", "reply to all", "forward", and "fw to contacts".

Figure 31. Incoming e-mail second or subsequent page

4.2.3 Edit Contact List

The contact list is edited using the **edit contacts** form (Figure 33). Because the structure of an individual contact group is the same as the structure of the top-level contact list, the same form is used to edit a contact group. The form is obtained by pressing the <edit contact list> button on the help page. It is also printed whenever a contact is added to the contact list, for example when the <add to contacts> button is pressed on an incoming e-mail to add the sender to the contact list.

Each entry on the contact list form shows the name on the contact, and the current e-mail folder in which e-mail from the contact is placed. Each entry has a checkbox which allows the contact to be selected, an <info> button which elicits a page of information about the user (Figure 44), and a <set mailbox> button which allows the e-mail folder associated with the contact to be changed via the **set e-mail folder** form (Figure 35).

The <delete selected> button at the bottom of the form deletes the selected contacts from the contact list (or contact group). The form is reprinted with the deleted contacts removed.

The <copy to group> and <move to group> buttons at the bottom of the form allows the selected contacts to be copied and moved to a particular group via the **copy contacts** form (Figure 34).

The contact list editing user interface flow is illustrated in Figure 32.

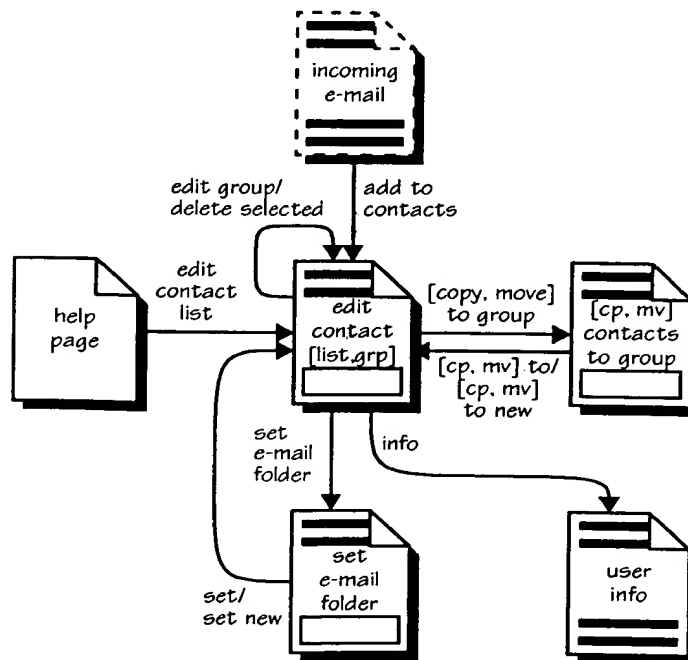


Figure 32. Edit contact [list,group] user interface flow

Edit Contact [List, Group] [page # of #]

group contact name

Select

<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name
<input type="checkbox"/>	info	contact name	set e-mail folder	e-mail folder name

delete selected copy to group move to group

Figure 33. Edit contacts page

The copy (move) contacts form lists all of the groups in the user's contact list, each with a <copy to> (<move to>) button which copies (moves) the selected contacts to the specified group. When a <copy to> (<move to>) button is pressed, an edit contacts form is printed which shows the updated membership of the destination group.

The copy (move) contacts form provides a <new group name> field with an associated <copy to new> (<move to new>) button. This allows a new group to be simultaneously created and selected as the destination for the copy or move.

[Copy, Move] Contacts to Group

[page 1 of #]

copy to new

new group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

copy to

group name

Figure 34. Copy contacts first page

The set e-mail folder form lists all of the folders in the user's mailbox, each with a <set> button which sets the folder as the contact's e-mail folder. When the <set> button is pressed, the edit contacts form is re-printed to reflect the new folder associated with the contact.

The set e-mail folder form provides a <new e-mail folder name> field with an associated <set new> button. This allows a new folder to be simultaneously created and selected as the contact's e-mail folder. The new folder has an associated set of checkboxes which allow the folder's e-mail delivery options to be specified.

Set E-mail Folder for
contact name

page 1 of #

set new new e-mail folder name

☐ print all e-mail
☐ print high-priority e-mail
☐ delete e-mail once printed

set Inbox

set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name
set	e-mail folder name

Figure 35. Set e-mail folder first page

4.2.4 List Mailbox

The contents of the user's mailbox are listed using the **mailbox** form (Figure 37) and the **e-mail folder** form (Figure 38). The mailbox form lists all of the user's e-mail folders, each with a set of checkboxes which reflect the folder's e-mail delivery options and allow the delivery options to be changed. Each folder entry also has a checkbox which allows it to be selected. The <count> next to each folder name indicates the number of un-printed (i.e. un-read) e-mails in the folder.

The <delete selected> button at the bottom of the form deletes all selected folders from the mailbox. The mailbox form is re-printed with the deleted folders removed.

The <update options> button at the bottom of the form applies any changes made via the checkboxes to the folders' delivery options.

Each folder entry also has <list> button which lists the contents of the folder via the e-mail folder form (Figure 38).

The mailbox listing user interface flow is illustrated in Figure 36.

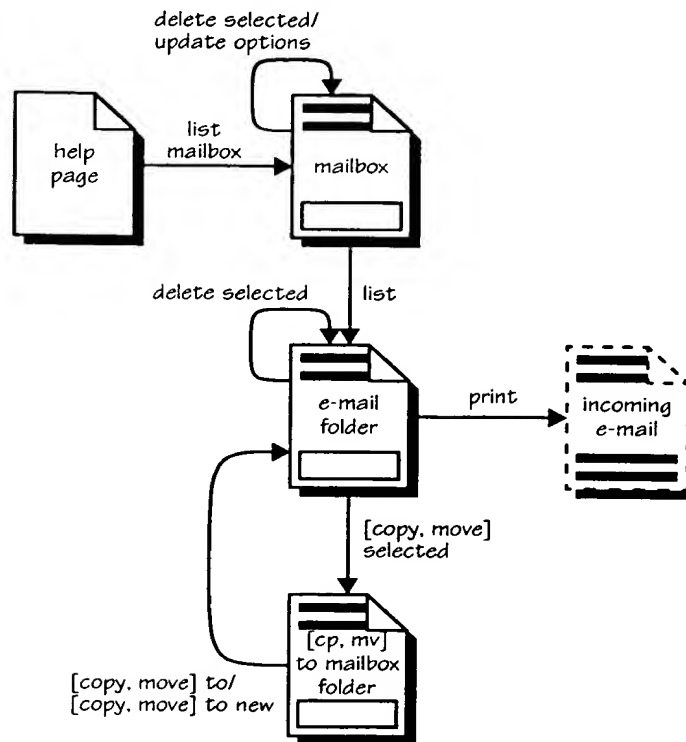


Figure 36. List mailbox user interface flow

The e-mail folder form lists all of the e-mails in the folder. Each single-line e-mail entry shows the sender of the e-mail, the subject of the e-mail, and the date and time the e-mail was received. The list can be sorted by sender, subject, or date and time. Each column which is not the current sort column has a <sort> button which allows the list to be sorted by that column. When a <sort> button is pressed the e-mail folder is sorted by the selected column and the form is re-printed.

[illegible]

Figure 37. Mailbox page

Where an e-mail is a reply to another e-mail, the subject of the replying e-mail is indented relative to the subject of the e-mail to which it is replying. Thus each thread of a discussion is clearly visible. Entries for e-mails which have never been printed are printed in bold to draw the user's attention.

Each e-mail entry has a checkbox which allows it to be selected, as well as a print button which allows the entire e-mail to be printed in the form of an Incoming E-mail (Figures 30 and 31).

The <delete selected> button at the bottom of the form deletes the currently selected e-mails. The e-mail folder form is reprinted with the deleted e-mails removed.

E-mail Folder

e-mail folder name

page # of #

Select	Sender	(sort)	Subject	(sort)	Recv'd	(sort)
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time
<input type="checkbox"/>	print		sender		subject	date & time

delete selected

copy selected

move selected

Figure 38. E-mail folder page

The <copy selected> and <move selected> buttons at the bottom of the e-mail folder form allow the currently selected e-mails to be copied or moved to a different e-mail folder via the copy e-mail form (Figure 39).

[Copy, Move] to E-mail Folder [page 1 of #]

☐ *print all e-mail*
☐ *print high-priority e-mail*
☐ *delete e-mail once printed*

Inbox

<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>
<input type="button" value="copy to"/>	<i>e-mail folder name</i>

Figure 39. Copy e-mail first page

The copy (move) e-mail form lists all of the folders in the user's mailbox, each with a <copy to> (<move to>) button which copies (moves) the selected e-mails to the specified folder. When a <copy to> (<move to>) button is pressed, an e-mail folder form is printed which shows the updated contents of the destination folder.

The copy (move) e-mail form provides a <new e-mail folder name> field with an associated <copy to new> (<move to new>) button. This allows a new folder to be simultaneously created and selected as the destination for the copy or move. The new folder has an associated set of checkboxes which allow the folder's e-mail delivery options to be specified.

4.2.5 Edit Barred User List

The barred user list is edited via the **edit barred users** form (Figure 41). The form is obtained by pressing the <edit barred user list> button on the help page. It is also printed whenever a user is barred, for example when the <bar> button is pressed on an incoming e-mail to add the sender to the barred user list.

The form lists all of the users in the user's barred user list. Each has a checkbox which allows it to be selected.

The <unbar selected> button at the bottom of the form deletes the selected users from the barred user list. The form is re-printed with the deleted users removed. When a user is deleted from the list that user is no longer barred.

The barred user editing user interface flow is illustrated in Figure 40.




Figure 40. Edit barred user list user interface flow

[illegible]

Figure 41. Edit barred users page

4.2.6 Global User Directory

The global user directory, since it is large, is navigated via a directory index (as described in Section 2 on page 6). It is obtained by pressing the <global user directory> button on the help page.

The **global user directory** page (Figure 43) contains an alphabetical list, by family name, of users in the name range covered by the page. Each entry has three associated buttons. The <info> button produces a **user information** page (Figure 44), the <e-mail> button generates an outgoing e-mail form (Figures 27 and 28) addressed to the corresponding user, and the <add to contacts> button adds the user to the user's contact list and prints the updated contact list (Figure 33).

The global user directory page also contains the standard index navigation buttons.

The global user directory user interface flow is illustrated in Figure 42. The index user interface flow and associated page layouts are as described in Section 2 on page 6.




Figure 42. Global user directory user interface flow

The user information page (Figure 44) contains the user's contact details as provided during user registration, subject to the user's privacy preferences. It also contains <e-mail> and <add to contacts> buttons which work in the usual way.

The <Internet e-mail address> is only included if the user information page is printed via the contact list for an Internet contact (as opposed to a Netpage contact).

[illegible]

Figure 43. Global user directory page

User Information

page 1 of 1

Name	[title, first name, initials, nickname, family name]
Address	[street address 1, street address 2, city, state, ZIP Code, country]
Home Phone	[home phone number]
Work Phone	[work phone number]
Mobile Phone	[mobile phone number]
Internet E-mail	[Internet e-mail address]

e-mail

add to
contacts

Figure 44. User information page

4.2.7 Business Card

A Netpage business card (Figure 46) serves a dual purpose. While it conveys the user's contact details in the usual way, it also acts as an e-mail authorization token. When the recipient of the card presses the <add to contacts> button on the card, the owner of the card is added to the recipient's contact list in the usual way, but the recipient is also added to the owner's contact list, authorizing the recipient to send e-mail to the card owner, should the card owner only accept e-mail from known contacts. The token only allows the recipient to be added to the owner's contact list once.

The card also contains an <e-mail> button which generates an outgoing e-mail form addressed to the card owner in the usual way.

The business card user interface flow is illustrated in Figure 45.




Figure 45. Business card user interface flow

[title] [first name] [initials] [nickname] [family name]
 [street address 1]
 [street address 2]
 [city] [state] [ZIP Code]
 [country]
 [work phone: work phone number]
 [mobile phone: mobile phone number]
 [home phone: home phone number]

e-mail

add to contacts

Figure 46. Business card

Since normal Netpage Printers don't support business card paper stock, Netpage business cards are provided by a service bureau on request. The **business card request** form (Figure 48) allows the user to specify which contact details to include on the business card, an optional background texture, and the number of cards required. Once the request is submitted, the cards are automatically printed by the bureau, delivered by mail to the user's address, and the small cost charged to the user's Netpage account.

The business card request user interface flow is illustrated in Figure 47.



Figure 47. Business card request user interface flow

Request Business Cards

page 1 of 1

Content		Background Texture
Title	<input type="checkbox"/>	<input type="checkbox"/> texture 1
First Name	<input type="checkbox"/>	<input type="checkbox"/> texture 2
Initials	<input type="checkbox"/>	<input type="checkbox"/> texture 3
Nickname	<input type="checkbox"/>	<input type="checkbox"/> texture 4
Family Name	<input type="checkbox"/>	<input type="checkbox"/> texture 5
Address	<input type="checkbox"/>	<input type="checkbox"/> texture 6
Home Phone	<input type="checkbox"/>	<input type="checkbox"/> texture 7
Work Phone	<input type="checkbox"/>	<input type="checkbox"/> texture 8
Photo	<input type="checkbox"/>	<input type="checkbox"/> texture 9

Quantity

50 ☐

100 ☐

200 ☐

500 ☐

Other

texture 10

texture 11

texture 12

texture 13

texture 14

texture 15

texture 16

texture 17

texture 18

texture 19

texture 20

texture 21

texture 22

texture 23

texture 24

Initialled by Administrator

Figure 48. Business card request page

4.2.8 Add Internet Contact

An Internet e-mail user can be added to a user's contact list using the **Internet contact registration** form (Figure 50). Since an Internet user is not otherwise known to the Netpage System, the form allows full name details to be specified, in addition to the Internet user's e-mail address.

The Internet contact registration user interface flow is illustrated in Figure 49.



Figure 49. Internet contact registration user interface flow

Add Internet Contact *page 1 of 1*

<i>Title</i>	<input type="text" value="title"/>	
<i>First Name</i>	<input type="text" value="first name"/>	<i>Initials</i> <input type="text" value="initials"/>
<i>Nickname</i>	<input type="text" value="nickname"/>	
<i>Family Name</i>	<input type="text" value="family name"/>	
<i>Internet e-mail</i>	<input type="text" value="Internet e-mail address"/>	

Figure 50. Internet contact registration page

5 Purchasing

The Netpage System provides efficient mechanisms for merchants to offer goods and services to customers, and for customers to select and pay for those goods and services.

A customer can order directly from a catalog browsed and printed dynamically via Netpage, as well as from a pre-printed catalog delivered by traditional means.

The customer's dynamic signature, in combination with the customer's uniquely identified pen, authenticates the purchasing transaction in a way which is impossible to forge. If the customer chooses to use the Secure Electronic Transaction (SET) payment mechanism built into the Netpage System, then the customer's payment card details are never revealed to the merchant, ensuring that card details can neither be accidentally nor maliciously compromised.

The purchasing mechanisms described in this section are generic, and apply to the purchasing of myriad kinds of goods and services.

5.1 PURCHASING OBJECT MODEL

The purchasing object model revolves around merchants and customers.

Each merchant has a globally unique identifier, as well as a name and other details. Each customer is related to a particular merchant, and has an identifier which is unique within the scope of the particular merchant. Each customer represents a particular Netpage user to the merchant. A Netpage user, on the other hand, can be the customer of any number of merchants.

Each merchant offers a range of products for sale, typically presented in the form of an organized catalog. Any Netpage user can browse the merchant's catalog, but the user must become an explicit customer of the merchant before being able to complete an order with the merchant.

The merchant class diagram is shown in Figure 51, while the customer class diagram is shown in Figure 52.

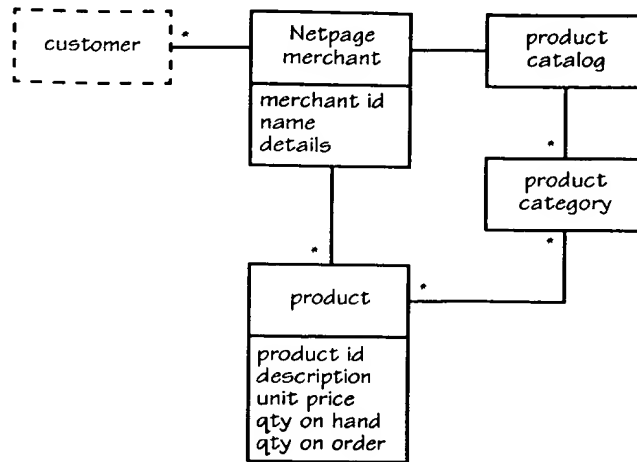


Figure 51. Merchant class diagram

The name and description of a customer derives from the details of the corresponding Netpage user.

Each customer has a history of payment methods and shipping addresses with the merchant. The most recent of each is recorded and used as the default on any new order.

Each customer has an account with the merchant which records the net balance of all invoices and received payments. Account payments, i.e. payments not associated with a particular invoice, are associated directly with the customer account.

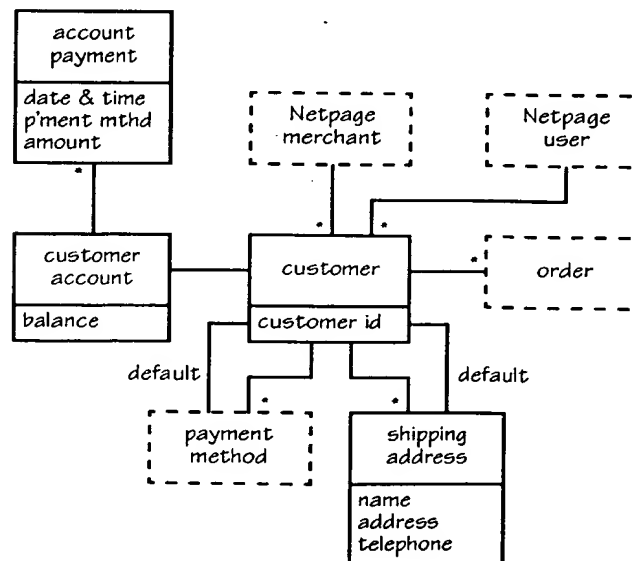


Figure 52. Customer class diagram

Each customer has a history of orders, each in a particular state of completion as reflected by its status. The order class diagram is shown in Figure .

Each order is uniquely identified within the scope of the merchant.

Each order has a status which indicates whether the order is pending (i.e. not yet submitted), active, partially or fully shipped, cancelled or completed. An order is not completed until full payment is received.

Each order consists of a number of order items, each of which specifies the quantity of a particular product offered by the merchant. Each order item also includes a backorder quantity and a quantity shipped. The backorder quantity is used to register a claim on future stock when there is insufficient current stock to fully satisfy the order item. The quantity shipped is used to record the cumulative quantity actually shipped.

Each order is associated with a particular payment method, shipping address, and shipping method.

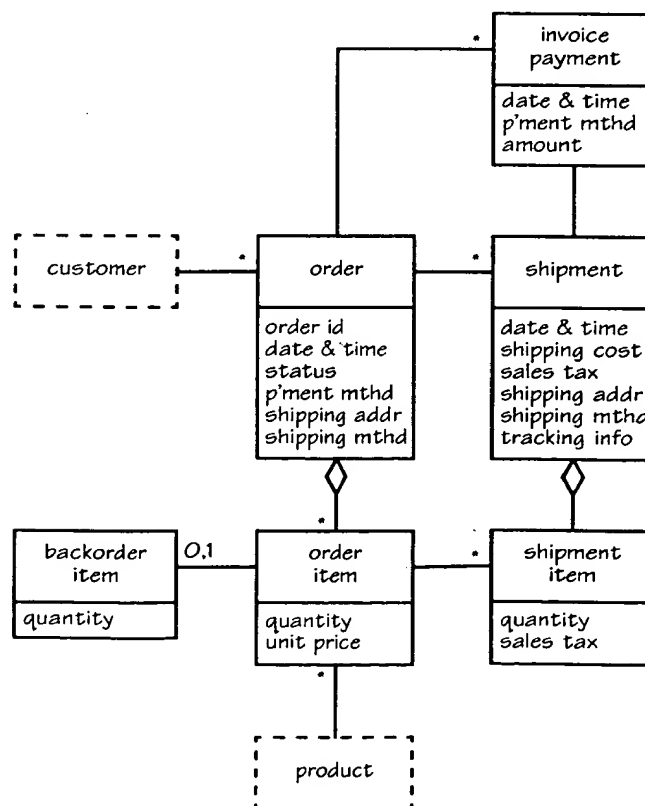


Figure 53. Order class diagram

The most common of the many possible payment methods are shown in the payment method class diagram in Figure 54. The set of payment methods supported by a particular merchant will be specific to that merchant, and may be different for different customers of the merchant depending on their credit ratings etc.

When the Netpage account payment method is used, the user's Netpage account is debited, and the transaction appears on the user's Netpage account statement.

When the customer account payment method is used, the user's customer account is debited, and the user is invoiced for payment.

When the pro forma invoice payment method is used, a pro forma invoice is issued requesting pre-payment of the order. When payment is received, the goods are shipped.

When the cash on delivery payment method is used, payment is expected when the goods are delivered.

When the SET payment card payment method is used, the corresponding SET payment card linked to the user is used to make payment according to the normal SET protocol.

When the payment card payment method is used, the specified payment card is used to make payment according to the normal protocol between the merchant and their acquiring bank or institution which handles the particular payment card type.

Invoice and pro forma invoice payments are normally made with a check or money order. Cash on delivery payments are normally made with a check or cash.

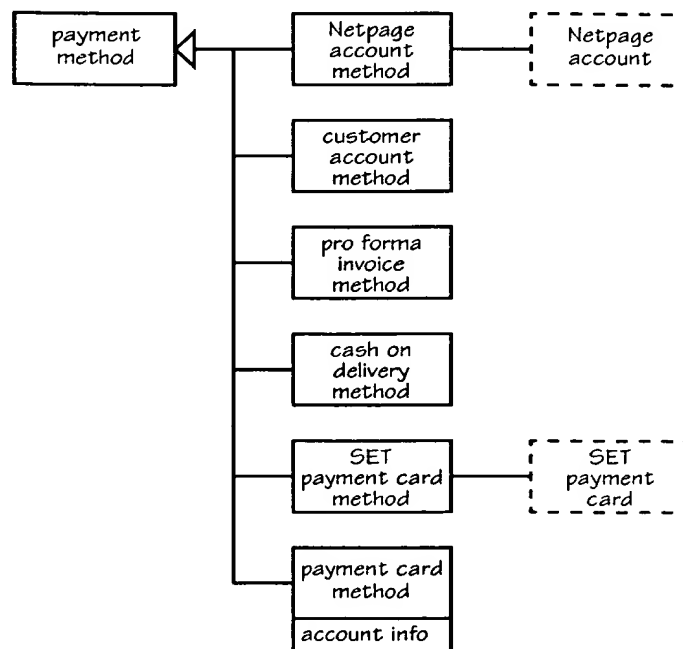


Figure 54. Payment method class diagram

The set of shipping methods which applies to a particular order depends on whether the order's shipping address is domestic or international with respect to the location of the merchant's warehouse. Typical domestic shipping methods include standard shipping, second day air, and next day air. Typical international shipping methods include standard shipping, airmail, and international priority.

The most common shipping methods are shown in the shipping method class diagram in Figure 55.

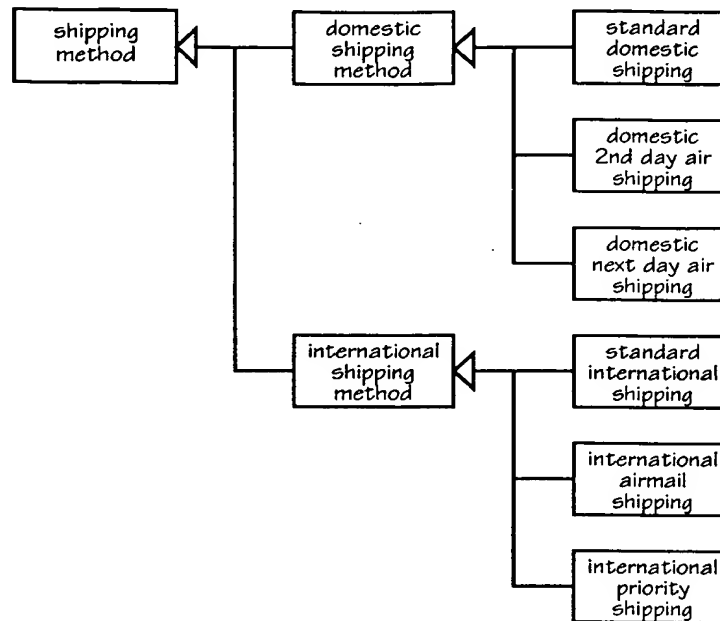


Figure 55. Shipping method class diagram

5.2 PURCHASING USER INTERFACE

The status of the order determines the allowable set of operations on the order. While the order is pending, the user can browse the merchant's catalog and add items to the shopping cart which represents the order. After the order is submitted, the user can still amend the payment method, shipping address and shipping method until the order is fully shipped. A change obviously only applies to items not yet shipped when the change is made. The user can delete order items or change the quantity ordered until the items are actually shipped. Quantity increases are subject to availability and payment approval.

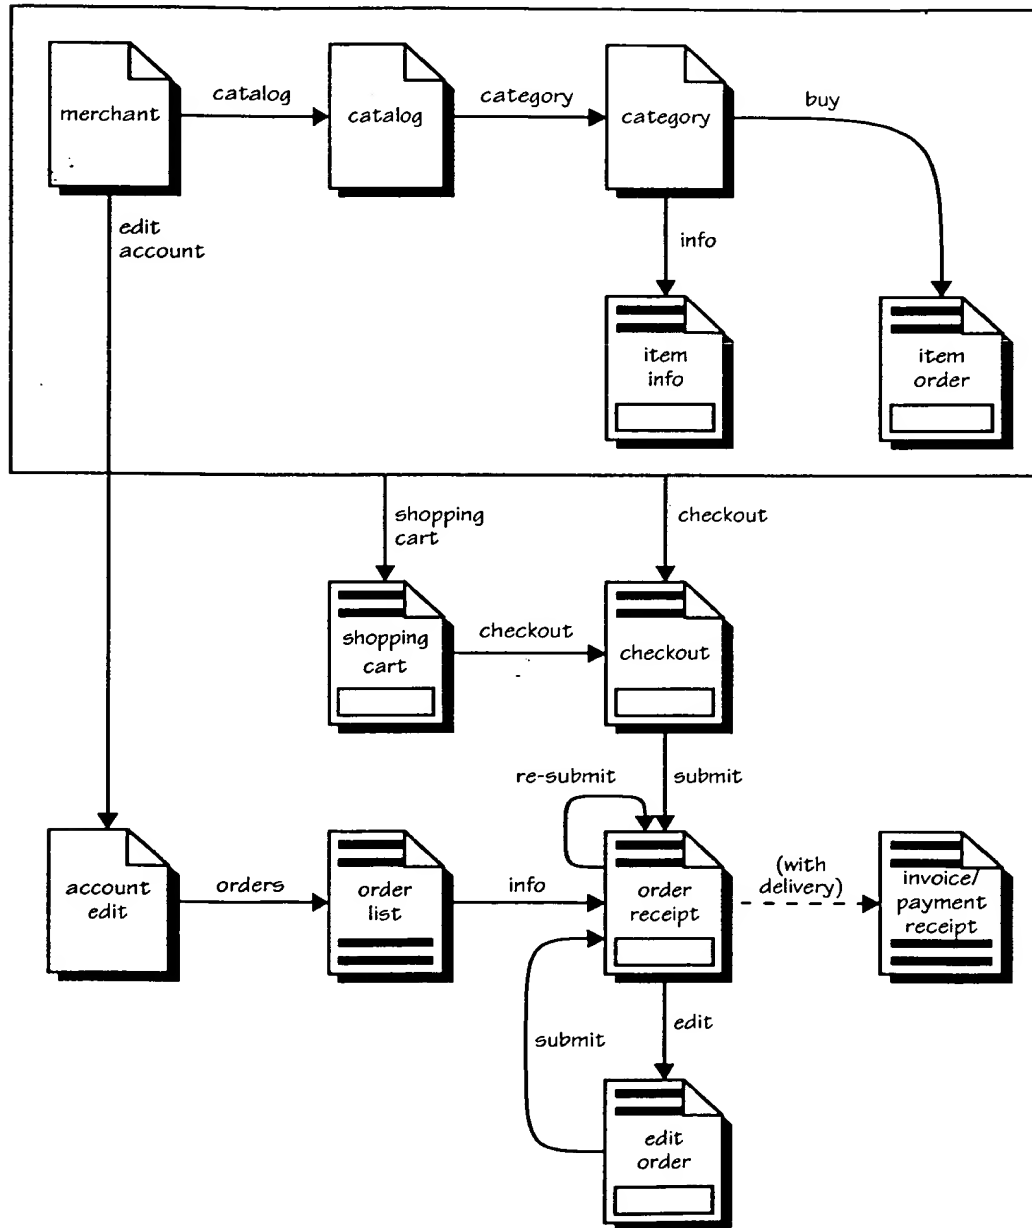


Figure 56. Online purchasing user interface flow

The overall purchasing user interface flow is illustrated in Figure 56.

5.2.1 Locate Merchant

A merchant can be located by browsing or searching the global Netpage directory or the user's own bookmarks. A merchant may also be encountered in the form of an advertisement in a Netpage publication. Both a directory entry and an advertisement typically link to the merchant's main page. This in turn links to the merchant's product catalog. If the user already has a copy of the merchant's main page from a previous encounter, then the catalog is immediately available.

5.2.2 Browse Catalog

A small product catalog is typically organized into a set of categories. The main catalog page simply lists the categories, and each category page lists all of the products in the category. This is indicated in the user interface flow in Figure 56.

If the product catalog is large, then it may still be organized into categories, but it may be impractical to print the entire contents of a category when a user wishes to browse the category. Instead the directory navigation techniques described in Section 2 can be used. A catalog search will often be specific to the kinds of products represented in the catalog.

5.2.3 Order Item

Many applications involve the user maintaining a collection of items. In the screen-based version of the application, the user's collection is typically shown every time an item is added or removed. The user therefore receives immediate feedback when the collection changes. The collection may reflect an e-mail recipient list, a shopping cart, etc.

In the Netpage version of the application, it is less practical to re-print the collection every time it changes. Item deletions and quantity changes can be marked up visibly, and so do not necessarily require immediate re-printing. Item insertions, however, are more problematic. In some cases the entire source list can be used as the input form, allowing it to be marked up visibly. This may be particularly apt if the source list already represents a user-specific subset of the complete source list, e.g. last week's shopping list rather than the grocery store's complete range. There will be cases, however, where the user will primarily be making selections from a complete catalog. The catalog may even be many hundreds of pages long, offset printed, and delivered in the mail. Since it is intended for multiple use, it is not suitable for visible markup.

One way to order via a large catalog is to print intermediate product-specific pages. Each catalog entry has both an <info> and an <buy> button. The <info> button elicits a full-page description of the product with a quantity field which, if filled in, adds the product to the shopping cart *and* provides a visible record of the item order. This avoids having to print the contents of the shopping cart every time an item is added to it. The <buy> button elicits a similar full-page description of the product but automatically adds it to the shopping cart. It contains a quantity field, already filled in with the quantity ordered by default, which can be used to vary the quantity ordered.

Depending on the nature of the order, it may be practical to also list the contents of the shopping cart on the product-specific page. However, it is the fact that the product page gives product-specific information that is of primary value to the user.

The product page is visibly linked to the order itself because it is tagged with the name of the current user, the order transaction number, and the order date and time.

5.2.4 Complete Order

The shopping cart represents the contents of the current order. The shopping cart is accessible from any merchant page, whether the page is part of the current order form or not. The shopping cart provides a convenient place to review the order and make final changes.

The checkout is also accessible from any merchant page. It provides a place to specify shipping and payment information, and finally to confirm the order.

The checkout form is broken up into five sections: shipping address, shipping method, payment method, order items, and order confirmation. Sections have default selections based on the user's most recent behavior.

The order confirmation section summarizes the order and captures the user's signature. It shows the current shipping address, shipping method, payment method, and order total.

Any change the user makes to the order on the checkout form is reflected on the form, since the change is made with the inking pen. It is not reflected, however, in the order summary in the order confirmation section. To produce a consistent checkout form, the user can simply press the <update> button at the bottom of the form after making a number of changes. An updated checkout form is immediately printed.

The <submit> button, when pressed, submits the order. Any handwritten change to any item information page, shopping cart page, or checkout page which is part of the order form is automatically interpreted and applied to the order. The order receipt, printed in response to the order submission, reflects the actual contents of the order.

The checkout form consists of a header (Figure 57), a section for selecting a previously used shipping address or specifying a new address (Figure 58), a section for selecting a shipping method (Figure 59), a section for selecting a previously used payment method or specifying a new payment card (Figure 60), a section for reviewing the list of order items and making last-minute changes (Figure 61), a section for signing the order to confirm payment (Figure 62), and a footer which contains the <update> and <submit> buttons (Figure 63).

The printed checkout form will vary in length depending on the number of previously used shipping addresses and payment cards, and the number of items in the order. In many cases it will fit conveniently on two pages printed on a single sheet of paper.

The order receipt lists the items ordered and summarizes the order parameters in much the same way as the order confirmation section of the checkout form (Figure 62).




Figure 57. Checkout page header

(1) Shipping Address

☐ shipping address

☐ shipping address

New Shipping Address

Name

Street Address

City

State ZIP Code

Country

Phone

Figure 58. Checkout page shipping address

(2) Shipping Method

<i>Domestic</i>	<i>International</i>
<input type="checkbox"/> Standard shipping	<input type="checkbox"/> Standard shipping
<input type="checkbox"/> Second day air	<input type="checkbox"/> Airmail
<input type="checkbox"/> Next day air	<input type="checkbox"/> International priority

Figure 59. Checkout page shipping method

(3) Payment Method

☐ payment card name & account expiry expiry date new expiry

☐ payment card name & account expiry expiry date new expiry

☐ Account

☐ Pro forma invoice

☐ Cash on delivery

New Payment Card

Name payment card name

Account payment card account

Expiry expiry date ☐ Visa ☐ MasterCard ☐ American Express ☐ Discover ☐ JCB

Figure 60. Checkout page payment method

(4) Items

	quantity	price	total price
item description	quantity		
item description	quantity		
item description	quantity		

Figure 61. Checkout page item list

(5) Confirm Order

<i>Subtotal</i>	subtotal
<i>Tax</i>	tax
<i>Freight</i>	freight
<i>Total</i>	total

Ship to

contact name, street address 1, street address 2,
city, state, ZIP Code, country, contact phone number

By

shipping method

Charge to

payment method

Signature required

signature

Figure 62. Checkout page order confirmation

update submit

Figure 63. Checkout page footer

5.2.5 Edit Account

Users can edit their customer details maintained by the merchant, including their list of shipping addresses and their list of payment methods. They can also print their order list,

and details of individual orders. They can edit any order which has not been fully shipped, as described earlier.

The order list gives the status for each order. Pressing the <info> button associated with each order yields a copy of the order receipt, but with an updated status for each order item. The status of an order item derives from the status of the order it is part, as well as the presence of shipment items and backorder items associated with the order item.

5.2.6 Edit Order

The order receipt allows the quantities of individual order items which are still active to be altered. Once alterations are made, the user must press the <submit> button. An updated receipt is produced.

To edit order parameters, the user can press the <edit> button on the order receipt. This produces an order editing form which is similar in structure to the checkout form. Once alterations are made, the user must press the <submit> button. An updated receipt is produced.

5.2.7 Receive Shipment

The shipment is accompanied by payment receipt, or, if the payment method is "cash on delivery", an invoice which allows the user to pay the delivery person by traditional means (e.g. cash, check, or payment card not processed through Netpage), or using a Netpage payment method listed on the form, simply by marking the appropriate payment method and signing the form with a Netpage pen.

A customer can therefore defer authorizing payment until the goods are delivered, even when using a payment method which allows the Netpage System to obtain immediate online payment authorization. Conversely, the delivery person can verify payment before handing over the goods.

6 Folder Hierarchy

6.1 FOLDER HIERARCHY

There are several instances where a Netpage user is able to organize a set of items into a hierarchy of folders. Examples include the user's contact list, favorites list, etc.

In each case the folder hierarchy consists of a root folder containing any number of folder entries. Each folder entry is either itself a folder containing further folder entries, or is a link to an actual folder item.

Folder entries may in general be members of more than one folder, although particular applications may not allow this. Both folders and folder items are typically named, and a folder is often presented to a user sorted by name, although sorting by other attributes can, of course, be equally valid.

When a folder item has an efficiently-retrievable name, and the folder is required to reflect the most up-to-date version of the name, then the folder item link is typically not named, and the name is instead retrieved from the folder item itself when required.

Each folder hierarchy is maintained by the Netpage Registration Server on the user's behalf, as part of the user information maintained by the server

The folder hierarchy class diagram is shown in Figure 64.

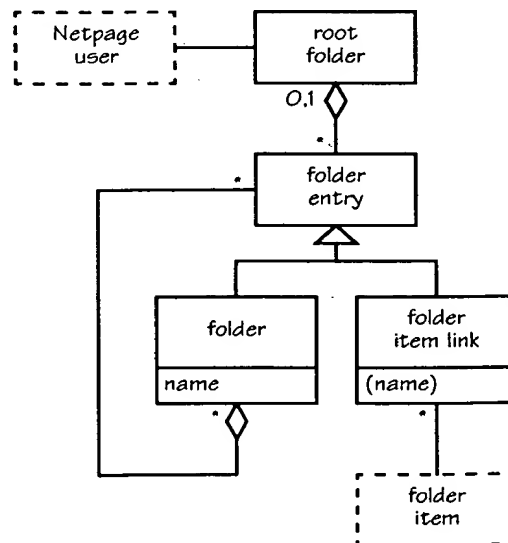


Figure 64. Folder hierarchy class diagram

7 Navigation Aids

7.1 FAVORITES LIST

Each Netpage user is allowed to maintain a list of “favorites” - links to useful documents on the Netpage Network. The list is maintained by the Netpage Registration Server on the user’s behalf, as part of the user information maintained by the server. The user can organize the list into a hierarchy of folders, edited and navigated in much the same way as the contact list described in Section 4.

The favorites list class diagram is shown in Figure 65.

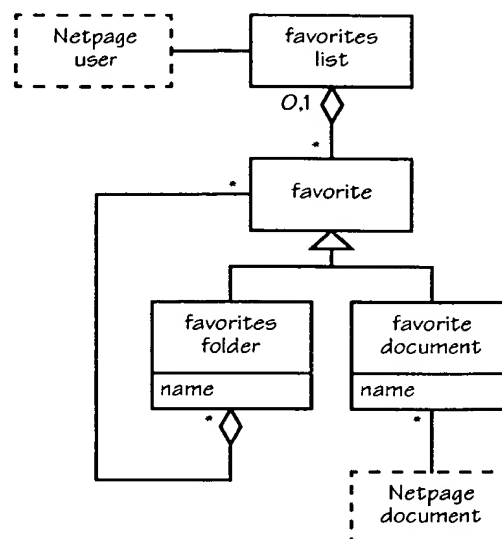


Figure 65. Favorites list class diagram

Each printed page of the favorites list contains three buttons for adding a favorite to the list: <add selected document>, <add selected page>, and <add selected link>. The user can select a region of a page in the usual way (i.e. by circumscription) and then press one of the three buttons. The <add selected document> button adds a link to the document which contains the selected page. The <add selected page> adds a link to the selected page. The <add selected link> button adds the first hyperlink contained in the current selection. This allows a document to be added to the favorites list without printing that document.

In each case the link is added to the root favorites list or a specific favorites folder depending on whether the page the button is pressed on contains a portion of the root list or a portion of a specific folder’s list. The help page also contains the three buttons for adding a favorite. These buttons always operate on the root list.

Each time a favorite is added a single double-sided page is printed showing the part of the list containing the new favorite.

Each entry in the favorites list is labelled with the name of the document. If the link is to an individual page, then the page number of that page is also shown.

If the user attempts to add a hyperlink which only makes sense in the context of its original location (e.g. a <submit> button on a form), then the operation is rejected with an appropriate error report.

7.2 HISTORY LIST

The Netpage Registration Server maintains a history list on each user's behalf, as part of the user information maintained by the server. The history list contains links to all documents accessed by the user through the Netpage System. The user can navigate the list by date and time, search the list by document name, delete entries from the list, and disable the server from maintaining the list.

The history list class diagram is shown in Figure 65.




Figure 66. History list class diagram

8 Drawing and Painting

The marking nib of a Netpage Pen is typically a ballpoint which produces a fixed-width line of a single fixed color. The Netpage System can nonetheless be used to draw and paint interactively using various line styles, colors, and brushes. Like in a graphical user interface drawing or painting application, the Netpage equivalent presents a palette of styles which can be used to draw or paint. The colors and styles used only become visible when the drawing or canvas page is reprinted.

Operation can be modal and modeless. When operating modally, the user simply selects colors and styles from the palette and draws or “paints” (by drawing) in the usual way on the surface of the page. There may be no indication of the current “mode” (i.e. color, pen, brush, etc.). The drawing or painting consists merely of fixed-width and fixed-color drawn lines until the page is reprinted. If the pen has an LCD then the LCD can indicate the current mode. However, there is no history of mode changes on the page itself.

When operating modelessly, color, pen and brush attributes are applied to a part of the picture after that part is drawn. Attributes can be applied in a modeless way by using a lasso.

8.1 LASSO

One way to support modeless pen attributes is to join elements of the drawing to pen attributes in a palette using a lasso drawn with the pen. The tail of the lasso is joined to the attribute in the palette, and the loop of the lasso circumscribes the part of the drawing to which the attribute is being applied.

The lasso is modeless since it provides a graphical representation of applying an attribute to a part of the drawing. Lassos can also be drawn with the pen in non-marking mode and a reprint can be requested (or can occur automatically) after each attribute is applied.

9 Viewing 3D Objects

When an application prints a view of a three-dimensional object on a Netpage, a number of controls can be provided to allow the user to change the view of the object. A set of controls is printed next to the image of the 3D object, and the user marks the controls and presses the <reprint> button to obtain a copy of the page containing the updated view.

The 3D object is projected to 2D according to a set of viewing parameters. These are normally defined in terms of a virtual camera, and typically include the following:

- *view point*, i.e. the point in 3D at which the camera is looking
- *camera position*, i.e. the point in 3D where the camera is located
- *camera twist*, i.e. the rotation of the camera about the viewing vector
- *view plane distance*, i.e. the distance from the camera position along the viewing vector to the plane onto which the 3D scene is projected
- *view plane window*, i.e. the 2D window on the view plane to which the scene is clipped
- *projection type*, i.e. parallel, perspective, etc.

The camera position is sometimes specified using intuitive parameters such as pan and tilt, where pan rotates the camera about the view point in the x-y plane, and tilt rotates the camera position about the view point in a plane which contains the camera position and the view point and is perpendicular to the x-y plane.

The 2D projected image of the 3D scene can also be subject to 2D viewing parameters, such as translation and scale, within the view plane window.

3D viewpoint controls can be implemented in many ways, including the following:

- *viewing direction*: a sphere surrounds a small copy of the object, with pan and tilt rings marked on the surface of the sphere (see Figure 67); the user specifies a new viewing direction by marking a new pan and/or tilt on the rings; the current pan and tilt settings are indicated on the rings, and the current viewing direction is indicated by a point on the surface of the sphere; the object can either be shown from a standard viewing direction or the current viewing direction.
- *viewing direction*: a small copy of the object is presented from opposite sides, with a transparent shaded hemisphere surrounding each copy; the user specifies a new viewing direction by marking a point on either hemisphere.
- *2D translation and scale*: the user circles an area of interest; the center of the circle specifies the translation, and the bounding box specifies the scale.
- *camera position*: a camera icon is shown in the position of the camera in the view of the object; the user specifies a new camera position by drawing a pen stroke from the camera icon to the desired position; the camera keeps pointing at the view point.
- *camera twist*: a handle is attached to the camera icon; the user specifies a new camera twist by drawing a pen stroke from the handle to the desired position.
- *automatically suggested views*: the object is shown in multiple automatically-generated views (e.g. left, right, top, bottom, and back views relative to the current camera parameters); the user simply selects one of the views.




Figure 67. 3D view control

Viewing parameters can also be specified via virtual sliders and numeric input. However, these forms of input are less intuitive and less user friendly, and so are more suited to professional applications where accuracy may be required.

Previously printed pages show a history of views. Unlike screen-based 3D navigation systems, the user cannot get lost because previous views of the scene are always available.

AUTHORING

10 Authoring Transition Strategies

10.1 OVERVIEW

During the early stages of Netpage Network deployment, publishers and merchants will be evaluating the costs and benefits of participating. In order to encourage such participation and grow the network, it is important to lower as far as possible any barriers to entry.

Although the Netpage protocols, servers and network infrastructure will be in place from the start, it will take time to develop generic high-level support for publishing, advertising and commerce. Early entrants may therefore be faced with significant development costs.

One way to mitigate this problem is to leverage development efforts occurring in the World Wide Web arena. The Netpage Network is a print medium which is highly compatible with existing print media. It is also an interactive online medium which is highly compatible with the Web. It is therefore well placed to exploit recent efforts to re-purpose print-oriented editorial content on the Web, and is also well placed to exploit Web applications defined at a sufficiently abstract level, as is now starting to become the norm.

Figure 68 illustrates these potential re-purposing relationships between print media, the Web, and the Netpage Network.




Figure 68. Re-purposing content and applications

10.2 RE-PURPOSING EDITORIAL CONTENT

Print publishers have been re-purposing content for the Web in various ad hoc ways almost since the inception of the Web. In more recent times, tools have begun to emerge which allow publishers to deploy content in multiple media from a single editorial database [4]. The key to such multi-purposing is the separation of content semantics from content layout.

XML is beginning to be widely used as a markup language to specify content semantics. During or after story creation, the journalist embeds semantic tags in the text. Once tagged, the story can be inserted into a page composition specific to a particular medium,

such as print or the Web, where it is laid out according to a medium-specific stylesheet cued to the semantic tags embedded in the text.

Several medium-independent editing and page composition systems are now emerging. Each supports output to both print and the Web, and can be extended in a modular fashion to support other output media such as Netpage. Atex' Omnex [2], System Integrators' Insight [10], and Worldweb.net's Expressroom [11,12] are notable examples. These systems integrate multiple-media publication into the core editorial activities of the publishing staff, rather than treating non-print media as separate and incidental.

10.3 RE-PURPOSING WEB APPLICATIONS

Web support for application development has improved enormously since the early days of raw HTTP and CGI. Application servers now provide high-performance back-end interfaces to database systems; transaction support with commit and rollback, and isolate the application user interface [1,6,7,8]. Very little of the effort involved in developing a modern Web application goes into Web-specific programming, and as a consequence the cost of re-purposing a Web application on the Netpage Network is not great. The functionality of the application remains unchanged, and only the detailed user interface flow and individual page designs need to change.

A number of turnkey and bureau Web commerce applications are emerging as well. iCat [5] is one of the better-known examples. It offers customized e-commerce hosting to any small company with a product catalog. By simply expanding the iCat system to treat the Netpage Network as an additional interface in parallel with the Web, all iCat-hosted businesses immediately gain entry to the Netpage Network with no effort on their part.

10.4 RE-PURPOSING WEB CONTENT

At a more mundane level, the Netpage Network provides an alternative interface to the Web itself, providing access to entire Web sites in printed form, with several Web pages printed on each physical page. Web pages appear in printed form yet remain fully navigable. The "brochure ware" which constitutes much of the Web is often more comfortably consumed in Netpage form.

Documents in Adobe PDF format, increasingly prevalent on the Web as companies strive for control over final print quality and shy away from re-coding information in inferior HTML, are even better candidates for Netpage delivery.

REFERENCES



11 Silverbrook References

12 Other References

- [1] Allaire, *ColdFusion*, <http://www.coldfusion.com/products/coldfusion/techinfo>
- [2] Atex, *Omnex*, <http://www.atex.com/products/omnex/omnex.htm>
- [3] Booch, G., J. Rumbaugh and I. Jacobsen, *The Unified Modeling Language User Guide*, Addison-Wesley 1999
- [4] Cavanagh, L., "Web Content Solutions for Newspapers Begin to Mature", *The Seybold report on Internet Publishing*, Volume 3, Number 11, July 1999
- [5] iCat, *iCat*, <http://www.icat.com>
- [6] Lotus, *Lotus Domino Application Server R5*,
<http://www.lotus.com/home.nsf/tabs/dominoapplicationserver>
- [7] Microsoft, *ASP*, <http://www.microsoft.com>
- [8] Oracle, *Oracle Application Server*, <http://www.oracle.com/appserver>
- [9] Rumbaugh, J., I. Jacobsen and G. Booch, *The Unified Modeling Language Reference Manual*, Addison-Wesley 1999
- [10] System Integrators Inc., *Insight*, <http://www.sii.com>
- [11] Votsch, V., "Worldweb.net spins Expressroom product out of Web service business", *The Seybold report on Internet Publishing*, Volume 3, Number 7, March 1999
- [12] Worldweb.net, *Expressroom*, <http://www.worldweb.net>

Netpage Extensions

draft version 0.1, 20 October 1999



Silverbrook Research Pty Ltd
393 Darling Street, Balmain
NSW 2041 Australia
Phone: +61 2 9818 6633
Fax: +61 2 9818 6711
Email: info@silverbrook.com.au

Confidential

Contents

NETPAGE EXTENSIONS	1
1 Overview	2
2 Netpage Printer	3
2.1 Electronic Ink Output	3
2.2 Braille Output	3
2.3 Printer Embedded in Digital Television Set	3
3 Netpage Pen / Sensor	4
3.1 Universal Drawing Implement	4
3.2 Output via LCD	4
3.3 Output via LED	4
3.4 Output via Tricolor LED	4
3.5 Interchangeable Nibs	5
3.6 Audio Output	5
3.7 Language Translation Output	5
REFERENCES	7
4 Silverbrook References	8
5 Other References	9

NETPAGE EXTENSIONS

1 Overview

This document describes a set of possible extensions to the core Netpage architecture. These extensions are described separately from the core architecture to ensure that the core architecture is clearly understood.

2 Netpage Printer

2.1 ELECTRONIC INK OUTPUT

Various “electronic ink” technologies are emerging, including E Ink’s electronic ink [2] and Xerox’ Gyricon electronic ink [3]. Both of these technologies involve bistable monochrome dots. They can be addressed externally, e.g. by passing them through a “printer”, or internally, by including addressing circuitry for each dot. There are several advantages to electronic ink. It has a display quality approaching that of ink on paper, since it is reflective rather than emissive or transmissive. It can be rewritten any number of times, just like a normal computer display, and arguably at fast enough rates to allow animation. And it is inexpensive enough to be bound into multi-sheet books, giving multi-sheet navigability just like paper books. In the short term it may only be practical to address it externally, and thus its main advantage may only be that it is easily reusable.

A reusable electronic ink page can be made Netpage addressable. Once “printed” onto an electronic ink page, that page can have all the functionality of a normal Netpage. What is required, however, is to introduce a level of indirection between the Netpage page id and the logical page, since the “content” of the physical page can change. Thus whenever the page is referenced by the pen, the tags are translated to the page id of the current Netpage displayed on the electronic ink page.

2.2 BRAILLE OUTPUT

A Netpage Printer can optionally contain a stamping mechanism for imprinting Braille characters on paper. The printer can thus produce an interactive Netpage which also contains the textual content of the page encoded in Braille. Documents delivered to the printer are personalized for the user in the usual way so that the text size and font match a useful Braille character size. Each sheet of paper is printed on one side only. Graphic elements on a page, such as images and buttons which do not contain text, are tagged with text in Braille. A Braille Netpage functions in the same way as a normal Netpage.

2.3 PRINTER EMBEDDED IN DIGITAL TELEVISION SET

A Netpage Printer can be embedded in a digital television set and receive data from the Netpage Network via a digital television “datacast” channel. It can use a low-bandwidth mobile telephone or pager channel to send data to the network. The printer thus leverages the communications capability of the digital television set, and provides a compelling use of a datacast channel. It also provides opportunities for seamlessly linking broadcast television content, e.g. infotainment, with printed material, as described in Section 9.2.3 in [1].

3 Netpage Pen / Sensor

3.1 UNIVERSAL DRAWING IMPLEMENT

A small Memjet printhead can be used to construct a universal drawing implement. A small cylindrical roller is in contact with the page. The roller spins freely about the axis of the stylus, so that as the roller is dragged across the page, it automatically orients itself so that it is at right angles to the direction of motion. The Memjet printhead is mounted behind the roller. It prints onto the back of the roller, and the roller transfers the printed image onto the page. A small cleaning station cleans any ink off the surface of the roller after it has contacted the page, so that the printhead always prints onto a clean surface. The printhead reproduces a contone color image via dithered bi-level CMY or CMYK. The stylus can be programmed to produce any colored, textured continuous line or paintstroke. The rate of printing is dictated by the speed of movement of the roller relative to the page. This can be determined in several ways, e.g. from the actual rotation of the roller, assuming reasonable friction between the roller and the surface, or by continuously imaging the surface and detecting movement in the same way as a second-generation optical mouse, or by sensing and decoding the map of a self-mapping surface.

3.2 OUTPUT VIA LCD

A Netpage Pen or similar sensing device can optionally incorporate a LCD which can be used to display modal, selection, an status information:

- mode name
- current color/texture for drawing/painting
- current line width for drawing
- current brush for painting
- first few words of current text selection
- image fragment of current image selection
- URI of hyperlink selection
- pen status (power, communication with printer, etc.)
- error messages

This allows immediate feedback without requiring an additional page to be printed. It potentially allows meaningful feedback when working offline.

3.3 OUTPUT VIA LED

A Netpage Pen or similar sensing device can optionally incorporate an LED which can be used to indicate that the pen is in a mode. The LED can server as a reminder that the user recently selected a non-standard mode, e.g. a colored brush for painting. An on-page button can be used to clear the current mode.

3.4 OUTPUT VIA TRICOLOR LED

A Netpage Pen or similar sensing device can optionally incorporate three LEDs - one red, one green, and one blue - which can be used in combination to indicate a continous-tone color currently associated with the sensing device, for example a virtual pen or brush color when the sensing device is being used to interact with a drawing or painting application.

3.5 INTERCHANGEABLE NIBS

A Netpage Pen or similar sensing and marking device can optionally support interchangeable nibs. Nibs can have different shapes, producing different line styles, and can have different colored ink supplies built in, producing different line colors. They can also produce different textures, depending on the composition of the nib. Each nib contains a code which uniquely identifies the type of the nib. The pen obtains the nib's identity code when the nib is attached to the pen. Any digital ink captured by the pen while the nib is attached is labelled with the nib's id.

Since some nibs will produce a different effect depending on their tilt and rotation, tilt and rotation must be embedded in the digital ink in such cases. Since the Netpage Pen is typically used in a somewhat fixed rotation, the nib rotation may be a function of how the nib is attached to the pen.

The nib id embedded in the digital ink can then be used by Netpage software to label the digital ink with line, color and texture characteristics which mimic those of the physical nib. A user can also choose to map a particular nib to an arbitrary pen or brush style, for example an air brush. Although the strokes reproduced by the Netpage software will no longer match the strokes produced by the physical nib, the user is able to distinguish strokes of different kinds on the physical page.

The nib id can also be used to signal the intent of the strokes produced with the nib. A particular nib might be designated as a text editing nib, for example, used solely to mark up text. By dedicating a nib to a particular function, it becomes easier for the software to disambiguate data (e.g. text) and commands (e.g. text editing commands).

Any number of Netpage Pens registered to the same user can be semi-permanently dedicated to particular drawing styles or functions by having function-specific nibs semi-permanently attached.

3.6 AUDIO OUTPUT

When a user interacts with a Netpage when out-of-range of a Netpage Printer, digital ink is buffered in the pen and is transmitted to and interpreted by the printer when the pen later comes within range of it. As a way of obtaining immediate feedback, pages can contain audio hyperlinks. When the user presses an audio hyperlink, the pen communicates with the user's mobile phone, the phone obtains the audio clip from the Netpage Network, and the phone plays the audio clip to the user. This assumes the pen and the phone share a short-range wireless technology such as Bluetooth, and that it is economical for the mobile phone to have an open data connection to the network such as promised by third-generation mobile telephony. It works best if the user is using the phone in hands-free mode together with a wired or wireless (Bluetooth) earphone.

The pen can also contain its own mobile telephony module and audio output, thus eliminating the need to interface with a mobile phone. Similarly, if audio output is restricted to only operate when the pen is within range of a printer, then only an audio output capability is required, either in the printer or in the pen.

3.7 LANGUAGE TRANSLATION OUTPUT

The Netpage System can offer an automated translation service. The user selects a word, a sentence, paragraph, or arbitrary block of text, and then presses a <translate> button, perhaps on the help page. The translation is displayed on the LCD, as described in Section

3.2. If the translation is too long to be displayed all at once, then it scrolls horizontally, slowly enough to be read.

Alternatively, the translation is delivered by audio output as described in Section 3.6.

REFERENCES

4 Silverbrook References

- [1] Silverbrook Research, *Netpage System Overview*, draft version 0.4, October 1999

5 Other References

- [2] E Ink, <http://www.eink.com>
- [3] Handley, K., "All the World's a Page",
http://www.ingersoll-rand.com/compair/oct_nov_97/octnovco/pag_1.htm



Self-Mapping Surfaces

Design Description

draft version 0.2, 20 October 1999



Silverbrook Research Pty Ltd
393 Darling Street, Balmain
NSW 2041 Australia
Phone: +61 2 9818 6633
Fax: +61 2 9818 6711
Email: info@silverbrook.com.au

Confidential

Contents

OVERVIEW	1
1 Self-Mapping Surfaces	2
2 Computer System User Interface	3
ARCHITECTURE	5
3 Tag Tiling	6
4 Tag Capture	10
APPLICATIONS	15
5 Digitizer	16
5.1 Large-Format Digitizer	16
5.2 Dynamic Handwriting Capture	16
5.3 Dynamic Signature Capture	16
5.4 Registered Projection	16
6 Situated User Interfaces	17
6.1 Device Control User Interface	17
6.2 User Interface on Toys	18
6.3 Defect Logging	18
7 Hybrid User Interfaces	19
7.1 Hybrid User Interface to Web Browser	19
7.2 Hybrid User Interface to Computer Application	19
HARDWARE	20
8 Tag Sensor	21
8.1 Wearable Sensor	21
REFERENCES	22
9 Silverbrook References	23
10 Other References	24

OVERVIEW

1 Self-Mapping Surfaces

A *self-mapping surface* is a physical surface in the real world which contains a perfectly registered map of itself which can be queried by an appropriate sensing device. A surface may be mapped sparsely or ubiquitously. The map may be encoded visibly or invisibly. The map may be defined at any required resolution. The map is defined in such a way that a local query on the mapped surface yields an unambiguous map reference both within the map and among different maps.

If a computer system contains information about features on a self-mapping surface, then such information can be efficiently retrieved based on map references supplied by the sensing device. The information thus retrieved can be used in many ways. It can be supplied directly to the operator of the sensing device in the form of audio, video, text or graphics, either projected onto the surface, virtually projected onto the surface by surface-registered display through the operator's head-mounted display, or displayed on a normal computer screen. It can, more generally, take the form of actions which are initiated by the computer system on behalf of the operator in response to the operator's interaction with the surface feature.

Self-mapping surfaces can be categorised according to whether the map is active or passive, i.e. whether the map actively participates when it is queried or not. Digitizing tablets, touch-sensitive screens and CRTs sensed by light pens are all examples of actively self-mapping surfaces. Active maps are typically only *partially* self-mapping, because they do not yield unambiguous inter-map references per se.

A passive map requires an active sensing device to perform the map query.

A surface map is encoded according to the following general principles:

- *the encoding consists of a tiling of tags* each of which encodes *surface id and position*
- tags are sufficiently small so that a single button click is guaranteed to sense at least one full tag
- tags are sufficiently large to be able to be applied reliably to the surface using available technology (e.g. printing), and to be able to be sensed reliably (e.g. optically)
- tags encode their position on the page to a higher precision than the tag density
- **tags need not tile a region with uniform density**
- higher-precision position information is deduced from the position of the tag in the sensed area
- each tag contains an *easily-recognized invariant structure (target)* which *aids initial detection*, and which *aids in undoing the effect of any warp induced by the surface or by the sensing process*
- the tag target is *rotationally asymmetric* so that the page can be imaged in any orientation and the orientation can be determined
- the tag target is *fault-tolerant* for immunity to surface damage and sensing errors
- the tag data is *fault-tolerant* for immunity to surface damage and sensing errors

An implementation of self-mapping surfaces is described in [1-5].

2 Computer System User Interface

A passive self-mapping surface can support a graphical user interface to a computer system, mediated by a sensing pen:

- *encoding of position to support context-dependent button clicking, drawing, writing and gesturing, or*
- encoding of *object id* to support *context-free* button clicking, or
- encoding of *page id* and position to support *context-free* button clicking, drawing, writing and gesturing
- *invisible* encoding to support clicking, drawing, writing and gesturing *anywhere, without visible interference with surface features*
- *fault-tolerant* encoding for *immunity to surface damage and sensing errors*

Minimal interface:

- consists of a physical page with a virtual page structure
- page surface contains features, e.g. graphical user interface elements
- features are typically printed onto the page surface
- physical page surface can be any surface, including the surface of a sheet of paper
- surface features are represented in corresponding page structure (there is a one-to-one mapping between the two)
- physical page is a (passive) self-mapping surface
- map uses fault-tolerant optical encoding
- encoding is typically printed onto the surface using invisible (e.g. IR) ink
- map is sensed by wireless optically sensing pen or stylus
- pen senses map when nib is close to or in contact with surface
- pen senses map reference at sufficient rate to capture human handwriting, i.e. at 200 dpi or greater and 100 Hz or faster [10]
- pen allows user to indicate "pen down" and "pen up" action, e.g. via button on shaft, or sprung button or pressure sensor coupled to nib
- pen generates digital ink with embedded map references
- digital ink consists of strokes
- stroke consists of timestamped positions bracketed by pen down and pen up
- pen transmits digital ink to base station
- page structure is known to base station
- base station interprets digital ink with respect to page structure, e.g. as operations on a graphical user interface

Page structure:

- page structure defines graphical elements and input elements
- graphical elements include text, graphics and images
- input elements include hyperlinks and form fields
- base station interprets digital ink as element selections, form field fill-ins, and hyperlink activations

Marking versus non-marking pen input:

- user gets visible feedback on input because pen has a marking nib
- pen also has a non-marking nib for command input (e.g. button presses)

Network:

- page structures are available on network
- pages are uniquely identified within particular domain
- base station retrieves page structures on demand by page id (from map reference)
- hyperlinked objects are available on network
- base station communicates with linked objects to handle hyperlink activations

Base station is printer:

- base station is a printer
- prints hyperlinked pages on demand
- prints subscribed pages autonomously

Persistence:

- page structures are automatically archived on network, making all published content persistent

Handwriting recognition:

- handwriting is automatically converted to text in text input field
- handwriting recognition uses dynamics in digital ink
- unrecognized handwriting can be automatically routed to human operator

Verified identity:

- pen acts as identity token
- signature is automatically verified in signature input field with respect to identity
- signature verification uses dynamics in digital ink, including pressure
- verified identity allows secure transaction authorization

ARCHITECTURE

3 Tag Tiling

Each tag is entirely self-sufficient. A tag can be located and decoded independently of surrounding information (e.g. other tags). Each tag identifies the surface on which it lies, and each tag specifies its own position on the surface. How a set of tags are arranged on a surface is somewhat immaterial, so long as the surface is tiled sufficiently densely in areas where tags are required, e.g. in the active areas of buttons or form fill-in fields. A tag tiling is sufficiently dense if the tag sensor can image at least one tag on a single "click" within any active area on the surface.

There are two reasons for using an *irregular* tiling:

- an irregular tiling is more easily mapped to a surface of arbitrary topography and topology
- an irregular tiling is less susceptible to damage from systematic artifacts; for example, if a tiling is *printed* onto a surface, a stuck nozzle can produce a column of dots which are always on or always off

An irregular tag tiling can be produced for a surface of arbitrary topography and topology as follows:

- (a) generate a Poisson-disk distribution of tag positions within an equilateral triangle of a certain size; ensure tag separation not just within the triangle, but also across the boundary between any two arbitrarily rotated adjacent copies of the triangle in any allowed triangle mesh
- (b) **compute a regular (quasi-equilateral) triangle mesh which approximates the surface;** a triangle mesh can approximate a surface of arbitrary topography and topology without introducing discontinuities or singularities, although surface detail which is finer than the size of the triangle may have to be ignored
- (c) map the tag triangle to each triangle of the mesh
- (d) project each mapped tag triangle onto the surface
- (e) render each tag in each projected triangle on the surface

This approach is based a pattern-based texturing approach described in [8]. That approach in turn relies on a surface re-tiling approach described in [11].

Figure 1 shows a regular tag triangle mesh which might be used to tag a planar surface such as a piece of paper. It also shows a distribution of eight tags within the tag triangle.

The process presumes that a polygonal model of the object exists and that rendering registered tags onto its surface is possible. A polygonal model of an actual object can be created from range-scanning data. A polygonal model may also exist because the object was manufactured from a polygonal model. The issue of rendering registered tags onto an arbitrary surface is not considered any further here.

A Poisson-disk distribution is a Poisson distribution which satisfies the Poisson-disk criterion [7], where the separation of any two samples is never less than a specified minimum. One way to generate a Poisson-disk distribution is to *jitter* the positions of regular grid points [6]. Another way is the so-called *dart-throwing* technique which generates one random position at a time and tests it against existing positions [7].




Figure 1. Tag triangle mesh and tag triangle

Each tag must now encode the following:

- the id of the tag position within the tag triangle
- the id of the tag triangle within the surface
- the id of the surface

A surface-specific description of the tag tiling can be retrieved based on the surface id, possibly from a network. The description includes the following:

- the shape and size of the tag triangle
- the distribution of tag positions within the tag triangle
- the arrangement of tag triangles on the surface

A particular tag tiling may be used for a class of surfaces, e.g. all sheets of Letter-sized paper, all double-donut Zap widgets, etc. If the tiling is regular in nature, e.g. if the surface is planar, then the triangle id can typically be translated into a surface position algorithmically rather than by table lookup.

Sensed tag information can be translated into a high-precision surface position as follows:

- translate tag triangle id into surface position
- translate tag position id into offset and add to surface position
- add offset of tag in imaged area to surface position

Tag tilings using non-triangular polygonal mesh units are also possible. Rectangular and hexagonal tilings of a plane are obviously possible, for example. Multiple different polygonal mesh units can also be used in a tiling, so long as they are identified in the tiling description. Multiple tag position distributions can also be used in a tiling, so long they are identified in the tiling description.

When a tag triangle is mapped onto a mesh triangle, distortion may occur because the mesh triangle is only *quasi*-equilateral and may vary slightly in size from the tag triangle. Furthermore, when the mesh triangle is projected onto the underlying surface, distortion may occur because of the greater detail of the surface. The first kind of distortion can be

reduced by computing tag distributions for a set of triangles of different shapes, and using the distribution for the shape which best fits the mesh triangle. The second kind of distortion can be reduced by using smaller tag triangles, although a tag triangle must always be large enough to support a proper Poisson-disk distribution. Alternatively, each triangle in the tag tiling description can be tagged with its distortion, so that the effects of the distortion can be included when a tag is interpreted.

However, if the surface features to which the tags are intended to relate are projected and rendered onto the surface using the same mesh, then the features will be distorted in the same way as the tags, and the features and tags will remain registered to each other.

No matter what approach is taken to handling distortion, it is important that the distortion does not cause the tag density to drop below the required threshold in any area, and conversely, that it does not cause tag overlap. If tags are found to overlap, then one of the offending tags can be removed, so long as this does not adversely affect tag density.

As stated above, when a Poisson-disk distribution of tag positions is generated for the tag triangle, tag separation must be enforced across the boundary between any two arbitrarily rotated adjacent copies of the triangle in any allowed triangle mesh. Figure 2 illustrates the number of unique triangle pair arrangements in an equilateral triangle mesh. For the purposes of ensuring tag separation there are six unique edge-to-edge arrangements. If the mirror image of the tag triangle is also used, then there are twenty-four unique arrangements. Vertex-to-vertex arrangements don't need to be considered because a tag will repel itself across an edge before it can repel itself across a vertex.




Figure 2. Number of unique triangle pair arrangements, without and with mirroring (figures in brackets)

During tag capture, the tag sensor must guarantee to capture at least one full tag when imaging any point on the surface. If tags are tiled on a regular grid then the tag sensor must image a circular area large enough to encompass the unit cell of the grid. The radius of the tag must also be added to the radius of the imaged area. If the tags are jittered on the grid, then the maximum jitter radius must also be added to the radius of the imaged area. This is illustrated for a triangular grid in Figure 3. A triangular grid provides the best tag packing of all regular grids.

If the center-to-center tag distance on a uniform triangular grid is k , the tag radius is r , and the jitter radius is j , then the radius m of the circular area which must be imaged by the sensor is given by:

$$m = \frac{k}{\sqrt{3}} + j + r \quad (\text{EQ 1})$$




Figure 3. Triangular tag tiling

4 Tag Capture

The sensor used to image and decode position tags on a self-mapping surface is typically in the shape of a pen or stylus.

There is wide variation in the tilt at which the pen is operated by different users and even by the same user at different times. This complicates the optics in the pen and the subsequent position tag decoding.

A reasonable tilt range needs to be supported, e.g. longitudinal tilt in the range 0° to 60° , and lateral tilt in the range -45° to $+45^\circ$. When the pen is vertical, i.e. normal to the paper surface, it has zero tilt. Figure 4 illustrates the lateral and longitudinal tilt planes.




Figure 4. Lateral and longitudinal pen tilt

Assuming simple fixed optics, tilt introduces two problems:

- out-of-focus blurring
- perspective distortion

The degree of blurring depends on the depth-of-field of the optics.

Perspective distortion makes position tag detection more difficult. Conceptually, the tilt must somehow be inferred to allow the perspective distortion to be removed before further image-processing takes place. Figure 5 illustrates tag perspective distortion.




Figure 5. Tag perspective distortion

Since it is impractical to image the area directly under the nib, tilt also affects the distance from the nib to the center of the imaged area. This distance must be known to allow an accurate nib position to be derived from the position determined from the position tag.




Figure 6. Tilt-induced effects

Figure 6 illustrates the effect of tilt on depth, perspective distortion and displacement in one dimension. D is the nominal distance from the lens plane to the paper surface when the pen is vertical; d is the actual tilt-induced distance to the center of the imaged area (B); a is the width of the imaged area; s is the distance from the pen axis to the imaging axis; θ is the pen tilt.

Three tilt-induced effects are of interest: the change in distance to the paper surface, the perspective distortion of the imaged area, and the change in the position of the center of the imaged area, relative to the origin (i.e. the pen nib position).

The tilt-induced distance to the paper is given by:

$$d = D - s(\cot \theta) \quad (\text{EQ 2})$$

If D is large compared with s , this effect is less significant. It needs to be quantified relative to the depth-of-field of the optics.

The perspective distortion of the imaged area, given by the ratio of imaged feature sizes at points A and C on the paper surface, is given by:

$$\frac{d-x}{d+x} \quad (\text{EQ 3})$$

Where:

$$x = \frac{a}{2}(\cot \theta) \quad (\text{EQ 4})$$

If D is large compared with a , this effect is less significant. It needs to be quantified relative both to the depth-of-field of the optics, and to the size of the position tag and the required image sensor resolution.

The change in position of the center of the imaged area (B) relative to the origin (O) is given by:

$$s(\sec \theta - 1) \quad (\text{EQ 5})$$

This effect is always significant for non-trivial tilt.

Naive tilt sensing based on gravity, e.g. using passive accelerometers, is inadequate because the page itself may be significantly tilted when the pen is being operated. Remember that the page is simply a piece of paper, not a tabletop electronic tablet.

Out-of-focus blurring can be eliminated by having variable-focus optics and an auto-focusing mechanism, e.g. a lens servo driven by image sharpness.

Tilt can be inferred in several ways. It can be computed from the perspective distortion observed of the otherwise rectangular position tag target areas. This is non-trivial because the targets have a complex structure which is potentially difficult to detect when they are distorted. This can be ameliorated by introducing additional structure into the overall tag tiling which is easier to detect than the tag targets themselves, although this further reduces the tag resolution. Tilt can also be computed from the distortion observed of a simpler pattern of light projected onto the page area being imaged.

The projected pattern must contain both longitudinal and lateral structure to allow both longitudinal and lateral tilt to be determined. The simplest pattern consist of four separated spots, two on the longitudinal axis and two on the lateral axis. Tilt along each axis is computed separately from the observed change in the corresponding inter-spot distance. This is illustrated in Figure 7.




Figure 7. Tilt-induced spot pattern perspective distortion
(arrows show direction of tilt,
dashed lines indicate changed inter-spot distance)

The capture of spot images and tag images can be time-multiplexed. The image sensor can easily be operated at twice the 100 Hz tag rate, and the (up to) 5 millisecond separation between spot and tag capture is unlikely to introduce error since tilt typically varies slowly over time.

The position tag macrodot pattern contains mostly high frequencies, and so reduces to a uniform gray if sufficiently low-pass filtered. If a tilt pattern spot is sufficiently large compared with the size of a macrodot, then it is easy to separate from the background of the position tag data. Once a spot is located, its center can be computed to greater accuracy.

As an alternative to projecting a spot pattern, each tag can include a structure which serves the same purpose: This structure should ideally project outside the tag body, to maximise its size and therefore maximise the accuracy with which perspective distortion is measured. Adjacent tags can have interleaved projections so long as they are non-overlapping, as illustrated in Figure 8.




Figure 8. Interleaved tags with external perspective targets

Once one spot is found, its partner is found along the same longitudinal or lateral axis. This makes detection more efficient. Similarly, since tilt typically changes slowly over time, the last known spot positions provide a good starting point for detection. At a more sophisticated level, longitudinal tilt and lateral tilt angular velocities can be maintained over time, allowing more accurate prediction of spot positions.

Perspective distortion can be eliminated entirely if the capture lens is always kept parallel to the page surface. Rather than achieving this using a moving lens, an alternative approach uses a wide-angle lens which always captures some part of the paper surface without distortion, as illustrated in Figure 9.




Figure 9. Undistorted imaging through wide-angle lens

The lens system can be further refined to provide different focal lengths at different positions, tuned to the expected distance to the page at the corresponding tilt. This may elimi-

nate the need for an auto-focus mechanism. However, since the same part of the lens will be imaging a significantly smaller depth when imaging the left side of the area of interest than when imaging the right side, the lens system may still provide insufficient depth-of-field to eliminate the need for variable focus.

The cost of this approach is that a larger area must be imaged, requiring a larger image sensor. It does not, however, place an additional burden on tilt decoding, since the tilt spots are restricted to a sub-area of the image area. Similarly, it does not place an additional burden on tag decoding, since once the tilt is known, the relevant tag sub-area is known.

The variable-focus wide-angle lens system can be designed using an automatic optimization technique such as a genetic algorithm.

APPLICATIONS

5 Digitizer

5.1 LARGE-FORMAT DIGITIZER

A passive self-mapping surface provides the basis for a large-format high-resolution digitizing tablet for a screen-based graphical user interface. The "digitizer" surface can be printed onto contact plastic which can then be glued to any surface, including the entire surface of a desk. The digitizer can be sold in sheets which can be used to tile a surface. The desk surface can itself be printed as a digitizer. The tags can be protected from wear by a tough overcoat or transparent laminate.

An imaging pen produces a stream of digital ink when interacting with the surface. The pen can be wired or wireless. A wireless pen may be able to communicate with a standard Bluetooth transceiver in the workstation in PC.

The user can map the large-format digitizer to an extended virtual workspace which exceeds the size of the workstation or PC screen. The system can be configured to ensure that the point of focus on the digitizer is always displayed on the screen, i.e. the screen can follow the pen.

5.2 DYNAMIC HANDWRITING CAPTURE

A passive self-mapping surface provides the basis for a low-cost high-resolution digitizing tablet for dynamic handwriting capture in conjunction with an imaging pen.

5.3 DYNAMIC SIGNATURE CAPTURE

A passive self-mapping surface provides the basis for a low-cost high-resolution digitizing tablet for dynamic handwriting capture in conjunction with an imaging and pressure-sensing pen.

5.4 REGISTERED PROJECTION

A self-mapping surface can be used as a large low-cost digitizing tablet for a screen-based graphical user interface system to a personal computer or workstation. If the screen output of the computer system is projected onto the self-mapping surface, in perfect registration with the surface map, then the surface becomes an interactive drawing surface, where drawing mediated by computer. It can thus be used for all kinds of interactive modelling, design, drawing, painting, etc.

The projector can include an imaging unit which senses the large-scale map of the surface (if it is encoded at multiple resolutions) and uses this information to control a servo which registers the projected image with the surface map.

Such a surface supports easy multi-user participation and can support the ultimate active desktop.

6 Situated User Interfaces

6.1 DEVICE CONTROL USER INTERFACE

There are many electrical and electronic devices in the home and office which have some kind of user interface for controlling their operation. These include office equipment, audio and video equipment, kitchen appliances, heating and cooling systems, etc. Even light sockets have on/off switches and sometimes dimmers, and in many countries power outlets have on/off switches.

Office equipment and entertainment equipment is increasingly networkable - i.e. devices come with interfaces to standard local-area networks, allowing their operation to be controlled over a network. The advent of standard wireless networking technologies such as Bluetooth may see an increase in networking support in such devices and appliances.

Some networking futurists predict that even light sockets will one day be networked, allowing their status to be queried and their function to be controlled remotely.

If a particular device is networked, and the network interface of the device allows the device to be fully controlled, then it is possible to completely separate the user interface to the device from the device itself. Furthermore, the user interface can then be implemented as a passive self-mapping surface, sensed by a wireless sensing device carried by the user. The sensing device must be small, unobtrusive and portable. It might be embedded in a ring, glove or finger glove worn by the user. It could also be embedded in custom eyewear with interaction controlled by gaze.

The self-mapping surface is coded with the unique id of the device, and is printed with the user interface of the device in the form of buttons etc. Being passive, the surface gives no feedback to the user as to the status of the device or its user interface. The base station for the sensing device carried by the user resolves operations on the self-mapping surface with reference to a user interface structure previously selected by the user from the device itself, or designed by the user using elements provided by the device.

User interfaces implemented in this way have several advantages:

- once the infrastructure is in place, the cost of manufacturing the user interface is removed from the cost of manufacturing each particular device; the cost of the user interface to a light socket, for example, is non-trivial
- a user interface can be replicated by simply printing another copy
- when a user interface becomes damaged it can be easily replaced
- when the firmware in a device is updated, a new user interface can be easily provided
- a user interface can be customized by arranging elements and printing the result
- a user interface with a particular style can be chosen from an available set
- multiple user-interfaces can be grouped together
- a user interface can be provided on the device itself by the same means
- access to individual user interfaces and functions can be controlled, since each sensing device is tied to a particular user's identity

6.2 USER INTERFACE ON TOYS

The surface of a toy can be printed inexpensively with the user interface to the toy as a self-mapping surface. An appropriate sensing device can be used to interact with the user interface, communicating by wire or wirelessly to either a personal computer, game console, or a processor embedded in the toy itself, to produce graphical output, sound effects and other behavior.

A toy sit-in submarine may have a large user interface containing dozens of submarine controls. The toy comes with a CD-ROM containing an associated interactive game which runs on a PC or game console. The game is driven by the self-mapping surface user interface contained in the submarine, perhaps mediated by a sensing glove which communicates wirelessly by Bluetooth with standard Bluetooth hardware in the PC or game console. The PC screen or console-connected television monitor provides the animated part of the user interface. The self-mapping surface provides the extended user interface.

Toys can be customized by selecting different user interfaces and printing them via a Netpage system [2]. Custom user interfaces can be created by assembling pre-defined user interface components.

6.3 DEFECT LOGGING

The entire surface of an object, e.g. a Boeing 747 aircraft, can be printed as a self-mapping surface. During maintenance of the object, an engineer can register the precisely location of detected defects or parts requiring maintenance, simply by touching the surface with an appropriate sensing device.

7 Hybrid User Interfaces

7.1 HYBRID USER INTERFACE TO WEB BROWSER

Self-mapping surfaces containing printed Web pages can be used to request linked Web pages for display in a screen-based Web browser. The printed Web pages provide an extended user interface to the Web. Printed pages in magazines can also be encoded as self-mapping surfaces so that their printed Web hyperlinks are also active.

A single page id can be used per page, or individual object ids can be used for each embedded hyperlink. The latter is more practical, since object ids just have to be translated to Web URLs, probably by some network server. In the former case, someone needs to keep track of the printed layout and logical content of each page. If the Netpage Web printing service is used [2], then this comes for free of course.

All that is needed is a special pen which talks to a Web browser rather than the printer, or a normal pen which can be placed in this mode.

7.2 HYBRID USER INTERFACE TO COMPUTER APPLICATION

The user interface to a GUI-based computer application running on a multi-tasking computer can be printed as a self-mapping surface on a user interface or command "card". The printed user interface can include a "digitizer pad" area for moving the GUI pointer relative to the application. Invoking any function of the application's user interface or moving the GUI pointer, automatically makes the application current - i.e. if the application is running in a Windowed GUI system then its window is brought to the front and made current. If the application is not currently running, then it is automatically launched.

The printed user interface for a text-oriented application can contain a printed keyboard, a general-purpose handwriting input text field, or both.

A personal computer system or workstation can thus potentially consist of a screen for displaying GUI output, a number of application-specific printed user interfaces, a sensing device (typically a stylus) for sensing user operations relative to the printed user interfaces, and a computer which receives wired or wireless transmissions from the sensing device, runs applications, and interprets sensed inputs relative to each application.

Each printed user interface "card" can be encoded with a unique page id specific to the application, and tagged with an attribute which instructs the personal computer or workstation to interpret operations on the page relative to a local instance of the application, even in a global networked Netpage environment.

If the computer is a network terminal connected to a LAN, an intranet, or the Internet, any interaction with the printed user interface launches or interacts with a networked instance of the application.

HARDWARE

8 Tag Sensor

8.1 WEARABLE SENSOR

A sensing device for a self-mapping surface worn on the body, e.g. in the shape of a ring, armband, bracelet, thimble, or mounted at the end of one of the fingers of a glove, etc., can be used to conveniently interact with a self-mapping surface. It can contain an active sensor for sensing the area of the self-mapping surface with which it is brought into contact or proximity, and transmits digital ink by wire or wirelessly to an associated base station, e.g. via Bluetooth.

REFERENCES

9 Silverbrook References

- [1] Silverbrook Research, *Netpage Applications*, draft version 0.2, October 1999
- [2] Silverbrook Research, *Netpage Extensions*, draft version 0.1, October 1999
- [3] Silverbrook Research, *Netpage Pen*, draft version 0.3, 7 September 1999
- [4] Silverbrook Research, *Netpage Printer*, draft version 0.3, 7 September 1999
- [5] Silverbrook Research, *Netpage System Overview*, draft version 0.4, October 1999

10 Other References

- [6] Cook, R.L., "Stochastic Sampling and Distributed Ray Tracing", in *An Introduction to Ray Tracing*, A.S. Glassner (Ed.), Academic Press 1989
- [7] Glassner, A.S., *Principles of Digital Image Synthesis, Volume One*, Morgan Kaufmann, pp.427-440
- [8] Neyret, F., and M-P. Cani, "Pattern-Based Texturing Revisited", *Computer Graphics Proceedings, Annual Conference Series, 1999*, pp.235-242
- [9] Plamondon, R. and G. Lorette, "Automatic Signature Verification and Writer Identification - the State of the Art", *Pattern Recognition*, Vol.22, No.2, 1989
- [10] Tappert, C., C.Y. Suen and T. Wakahara, "The State of the Art in On-Line Handwriting Recognition:", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol.12, No.8, August 1990
- [11] Turk, G., "Re-Tiling Polygonal Surfaces, *Computer Graphics*, Vol.26, No.2, July 1992, pp.55-64